Special Interest Group on Design Automation
ACM/SIGDA E-NEWSLETTER, Vol. 55, No. 10

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. **US-UK Tech Prosperity Deal Brings Stability and Momentum**
The United States and the United Kingdom have signed the Technology Prosperity Deal (TPD), an ambitious agreement spanning artificial intelligence, nuclear energy, and quantum computing. President Donald J. Trump and Prime Minister Keir Starmer's move represents a significant boost to transatlantic cooperation at a time marked by growing divisions in global technology.

2. **Europe Drives to Dominate Photonics**
From interconnect to optical computing and even quantum computers, photonics in Europe is seeing a significant boost this month. By the end of the decade, the global photonics market is expected to exceed €1.5 trillion, driven by the increasing need to increase the bandwidth of data center AI chips while reducing the power consumption.

3. **Groq Secures $750M to Scale AI Inference Infrastructure**
Groq has raised $750M in new funding at a $6.9 billion valuation, fueling its ambitions to expand AI inference infrastructure globally. The financing highlights how investors are betting big on inference acceleration that can keep pace with surging demand for AI workloads.

4. **AI Surge Creates Nearline HDD Shortages, SSDs Poised to Benefit**
New research by TrendForce indicates that inference AI applications are generating massive data volumes, straining the global storage ecosystem. Nearline HDDs are now in short supply, with lead times ranging from weeks to over a year. This shortage is forcing cloud service providers (CSPs) to reassess their tiered storage architectures.

5. **Nvidia Deal Paves the Way for Intel Buy**
A key deal announced today paves the way for Nvidia to buy the design business of struggling chip maker Intel. The deal is significant in many ways. Nvidia will work with Intel to create a custom x86 processor with the NVLink interface for the

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 Event-based Vision, 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

AI datacenter market, and Nvidia will also supply GPU chiplets to Intel for consumer AI PCs.

6. **Quantum Motion Unveils First Silicon CMOS Quantum Computer**
Quantum Motion has launched the industry's first full-stack quantum computer built on standard silicon CMOS technology. The company installed the system at the UK's National Quantum Computing Centre (NQCC), marking a milestone for silicon-based approaches to quantum hardware.

7. **Cadence Adds NVIDIA DGX SuperPOD Digital Twin for AI Data Center**
Cadence has expanded its Reality Digital Twin Platform library with a new digital twin model of the NVIDIA DGX SuperPOD with DGX GB200 systems, giving data center designers the ability to simulate next-generation AI infrastructure before committing to physical deployments. The move strengthens Cadence's positioning in the fast-growing AI factory market, where design accuracy, power efficiency, and rapid scaling are critical.

8. **Mitsubishi Electric Launches Compact DIPIPM Power Semiconductor Modules**
Mitsubishi Electric has introduced a new compact version of its DIPIPM (Dual Inline Package Intelligent Power Module) power semiconductor modules, designed for use in consumer and industrial equipment such as packaged air conditioners and heat pump systems. The new Compact DIPIPM lineup includes the PSS30SF1F6 (30A / 600V) and the PSS50SF1F6 (50A / 600V), with sample shipments starting September 22.

9. **SiFive New RISC-V AI IP with Scalar, Vector, and Matrix Compute**
SiFive has expanded its RISC-V AI portfolio with the launch of its 2nd Generation Intelligence family, introducing five new processor IPs designed to accelerate AI workloads from the far edge to the data center. According to SiFive, the additions include two new cores — the X160 Gen 2 and X180 Gen 2 — alongside refreshed versions of the X280, X390, and XM.

10. **Quantinuum Raises $600 Million to Drive Scalable Quantum Computing**
Honeywell has announced a $600 million equity capital raise for its quantum computing arm, Quantinuum, at a pre-money valuation of $10 billion. This latest funding round marks a major step in advancing high-performance quantum computing toward practical, large-scale applications.

SIGDA E-News Editorial Board

SIGDA Awards

1. Best Paper Award @ MLCAD Symposium 2025
https://mlcad.org/symposium/2025/

ORFS-agent: Tool-Using Agents for Chip Design Optimization
Amur Ghose, Andrew Kahng, Sayak Kundu, Zhiang Wang
University of California San Diego, USA

What is Event-Based Vision?

Contributing author: Charlotte Frenkel <c.frenkel@tudelft.nl>
Assistant Professor at the Microelectronics Department, Delft University of Technology, Delft, The Netherlands
AE: Alberto Marchisio <alberto.marchisio@nyu.edu>

Event-based cameras, instead of capturing full images at a fixed frequency of a few tens of Hz, generate events on a per-pixel basis, representing local changes in temporal contrast exceeding a given threshold [1]. This leads to several interesting properties, such as increased sparsity as static background information does not generate threshold-crossing events, or fine-grained temporal resolutions down to a few microseconds. Event-based vision thus outlines new opportunities for low-power and/or latency-critical applications, such as in the automotive domain, for tiny flying drones, or for smart glasses.

However, the field of event-based vision has not yet reached the point of demonstrating microsecond-range latency, real-time processing of high-resolution data. This stems from the fact that event stream data is highly irregular, in stark contrast with the regularity of fixed-interval full frames that facilitates data reuse and parallelism in downstream processing hardware. Event-based vision thus needs novel algorithmic and hardware techniques [2].

Accumulating events within a time window (a process also referred to as binning, or the dense-frame approach), thereby yielding frame-like data, is a good intermediate step [3]; it allows reusing conventional vision algorithms with an adjustable temporal resolution based on the selected time window. As the resolution reduces, the number of accumulated events gets smaller. The resulting frames can still be quite sparse, making this approach efficient if downstream hardware can exploit sparsity in input data [4]. Spiking neural networks offer a relevant example, not only for sparsity-awareness, but also for their ability to handle temporal information via stateful integrate-and-fire neurons [5]. Efficiently handling the memory footprint of this extra state is the subject of ongoing research, for example, with the extension of in-memory computing hardware architectures to event-based, sparsity-aware scenarios [6,7].

In order to truly harness the microsecond-range latency opportunities offered by event-based cameras, fully event-driven approaches are the equivalent of pushing binning approaches to the extreme, i.e., down to the sensor's temporal resolution. As doing so forgoes the regularity induced by binning, it is necessary to find new parallelism and data reuse opportunities to compensate for the irregular nature of the input. Graph neural networks (GNNs) emerge as a promising approach, where events are represented as nodes of a sparse graph. Indeed, novel techniques [8,9] have introduced event-driven processing to GNNs, where the graph is dynamically expanded with new nodes as new events are received, and where information is exchanged locally among nodes that are within a given spatiotemporal distance. Event-based GNNs open new hardware/algorithm co-design opportunities, for example, via the parallel execution of multiple GNN layers, thereby compensating for the irregularity of

Paper Deadlines

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

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

FCCM'26 – IEEE International Symposium On Field-Programmable Custom Computing Machines
Atlanta, Georgia, USA
Abstracts due: Jan. 10, 2026
Deadline: Jan. 17, 2026
May 13-16, 2026
https://www.fccm.org/

fully event-driven execution and enabling latencies as low as 16µs per event with low-cost edge FPGA platforms [10].

Down the line, one does not need to oppose local, microsecond-latency, and event-driven to global, millisecond-latency, and frame-based. On the contrary, these two views complement each other toward vision systems that thrive in latency-critical application scenarios with high performance requirements [11].

References:

[1] P. Lichtsteiner, C. Posch and T. Delbruck, "A 128 × 128 120 dB 15 µs latency asynchronous temporal contrast vision sensor," IEEE Journal of Solid-State Circuits, vol. 43, no. 2, pp. 566–576, Feb. 2008.

[2] G. Gallego, T. Delbruck, G. Orchard, C. Bartolozzi, B. T. Meyer, A. Mueggler, M. Liu, S. Leutenegger, and D. Scaramuzza, "Event-based Vision: A Survey," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 44, no. 1, pp. 154–180, Jan. 2022.

[3] M. Ji, Z. Wang, R. Yan, Q. Liu, S. Xu, and H. Tang, "SCTN: Event-based object tracking with energy-efficient deep convolutional spiking neural networks," Frontiers in Neuroscience, vol. 17, no. 1123698, Feb. 2023.

[4] M. Shi, A. Kneip, N. Chauvaux, J. Sun, C. Frenkel, and M. Verhelst, "Sparsity-aware hardware: From overheads to performance benefits," IEEE Solid-State Circuits Magazine, vol. 17, no. 2, pp. 61-71, June 2025.

[5] K. K. Eshraghian, M. Ward, E. O. Neftci, X. Wang, G. Lenz, G. Dwivedi, M. Bennamoun, D. S. Jeong, and W. D. Lu "Training spiking neural networks using lessons from deep learning," Proceedings of the IEEE, vol. 111, no. 9, pp. 1016-1054, Sep. 2023.

[6] A. Agrawal, M. Ali, M. Koo, N. Rathi, A. Jaiswal, and K. Roy, "IMPULSE: A 65-nm digital compute-in-memory macro with fused weights and membrane potential for spike-based sequential learning tasks," IEEE Solid-State Circuits Letters, vol. 4, pp. 137–140, June 2021.

[7] N. Chauvaux, A. Kneip, C. Posch, K. Makinwa, and C. Frenkel, "An event-based digital compute-in-memory accelerator with flexible operand resolution and layer-wise weight/output stationarity," Proc. of IEEE International Symposium on Circuits and Systems (ISCAS), May 2025.

[8] S. Schaefer, D. Gehrig, and D. Scaramuzza, "AEGNN: Asynchronous event-based graph neural networks," Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 12371-12381, June 2022.

[9] T. Dalgaty, T. Mesquida, D. Joubert, A. Sironi, P. Vivet, and C. Posch, "HUGNet: Hemi-spherical update graph neural network applied to low-latency event-based optical flow," Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3953-3962, June 2023.

[10] Y. Yang, A. Kneip, and C. Frenkel, "EvGNN: An event-driven graph neural network accelerator for edge vision," IEEE Transactions on Circuits and Systems for Artificial Intelligence, vol. 2, no. 1, pp. 37-50, Mar. 2025.

[11] D. Gehrig and Davide Scaramuzza, "Low-latency automotive vision with event cameras," Nature, vol. 629, no. 8014, pp. 1034-1040, May 2024.

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

Upcoming Conferences

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

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

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

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.

EDITOR'S PICK FROM ACM TODAES ISSUE 4, 2025

HAPE: Hardware-Aware LLM Pruning For Efficient On-Device Inference Optimization

Wenqian Zhao, Lancheng Zou, Zixiao Wang, Xufeng Yao, Bei Yu
The Chinese University of Hong Kong, Hong Kong

Abstract: This paper introduces a practical, post-training pruning process to address the computational challenges of deploying LLMs on diverse edge

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/

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

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

ISPD'26 – ACM Int'l Symposium on Physical Design
Bonn, Germany
Mar 15-18, 2026

devices. Rather than focusing solely on sparsity, the proposed hardware-aware framework is designed to maximize actual inference speed-up by evaluating the sensitivity of different pruning structures. This method breaks from conventional layer-wise pruning by fusing layers to enable cross-layer optimization and integrate compilation passes to ensure sparsity patterns are executable for real acceleration. This approach eliminates the need for resource-intensive retraining while providing tailored performance and accuracy trade-offs for target hardware.

Continue reading more on ACM DL.

Technical Activities

1. **Project Aims To Make Real-Time Digital Twin Of The World**

A group of partners, led by geospatial simulation technology company Aechelon, has launched Project Orbion, a new initiative to create a real-time digital twin of the world. This digital twin will be based on real-time satellite imagery, radar, video photogrammetry and AI-created synthetic data, and is intended for commercial use cases…

2. **Microchip Thermocouple Conditioning IC Touts ±1.5°C System Accuracy**

Microchip Technology Inc.'s single-chip, four-channel MCP9604 integrated thermocouple conditioning IC delivers up to ±1.5°C accuracy and provides an alternative to discrete and multichip thermocouple conditioning solutions that can introduce errors and add system design complexity. The device is designed for critical temperature measurement systems used in production-line applications ranging from chemical and food processing, manufacturing process control and medical and HVAC equipment to refrigerated, cryogenic and other carefully controlled environments…

3. **Spinning Rust Hits a Wall: SSDs Rewrite the Storage Rules for AI Workloads**

AI is no longer asking for a storage upgrade—it's rewriting the contract. The moment machine-learning jobs shifted from lab demos to 24×7 production, the hard disk's spinning platters became the bottleneck. What once won budgets on dollars-per-terabyte now loses on time-per-epoch: every seek, every rotational delay is a GPU starved. Hyperscalers, national labs, and supercomputing centers aren't gradually favoring flash—they're sprinting to it, because training the next trillion-parameter model can't wait for a mechanical arm to swing…

4. **SK hynix Maintains Memory Leadership with First HBM4**

SK hynix has completed the development of HBM4 memory chips, which are now ready for mass production, six months after the Korean chipmaker delivered samples of its 12-layer HBM4 to key customers, including Nvidia. SK hynix calls the launch of the first HBM4 device a symbolic turning point beyond artificial intelligence (AI) infrastructure limitations…

http://www.ispd.cc/

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

DATE'26 - Design Automation and Test in Europe
Verona, Italy
April 20-22, 2026
http://www.date-conference.com

Job Positions

University of Texas Austin, US

Job Title: Assistant Professor of Computer Engineering

Description: The Chandra Family Department of Electrical and Computer Engineering at The University of Texas at Austin has multiple faculty openings with a start date of Fall 2026 for tenure-track faculty positions. Candidates in all areas of electrical and computer engineering will be considered, especially as they align with the strategic focus of the department in: AI/ML, including generative AI foundations and applications; AI for systems and systems for AI; Computer architecture, with a focus on heterogeneous computing systems; Signal processing, data science, communications and information theory; Software engineering, with a focus on AI-managed middleware, software generation and applications. We are particularly interested in candidates who will contribute to supporting our student body through their teaching, research, and service. Submission by October 15, 2025, is strongly encouraged. For more information, please refer to https://facultyvacancies.com/assistant-professor-of-computer-engineering,i43179.html.

University of Toronto, Canada

Job Title: Assistant Professor of Computer Science

Description: The Department of Computer Science in the Faculty of Arts and Science at the University of Toronto invites applications for two full-time teaching stream positions in the area of Computer Science. The appointments will be at the rank of Assistant Professor, Teaching Stream with an anticipated start date of July 1, 2026. Applicants must have earned a Ph.D. degree in Computer Science or a related field by the time of appointment or shortly thereafter. Alternatively, candidates are required to have (i) a Masters in Computer Science or related field with (ii) at least 18 months of excellent teaching experience in a degree-granting program/post-secondary institution, and (iii) demonstrated excellent scholarly or creative professional activity in areas such as, but not limited to, exemplary Computer Science teaching practices, development of pedagogical software tools, course or curriculum development, or engagement with Computer Science Education research. For more information, please refer to https://facultyvacancies.com/assistant-professor-of-computer-science,i43127.html.

Indiana University Bloomington, US

Job Title: Assistant/Associate/Full Professor of Computer Science

Description: The Indiana University Luddy School of Informatics, Computing and Engineering at IU Indianapolis invites applications for multiple open rank tenured or tenure-track assistant, associate or full professor positions in computer

science. Appointment start date is ideally January 1, 2026 or sooner. Candidates must demonstrate an outstanding scholarly record of research, exhibited by high-impact peer-reviewed publications and a forward-looking, externally funded research agenda as Principal Investigator (PI). Research expertise in the following areas of computer science will be considered: Fundamentals of Artificial Intelligence and Machine Learning, Robotics, Computer Vision, Cyber security, Cyberphysical Systems, and Quantum Computing. Candidates with experience in applications to biomedical and health sciences are strongly encouraged to apply. As the state's premier urban research institution, IU Indianapolis is committed to being a welcoming and inclusive campus community. We seek candidates who will pursue the highest standards of academic excellence and whose research, teaching, and community engagement efforts contribute to welcoming, respectful, and inclusive learning and working environments for our students, staff, and faculty. For more information, please refer to https://facultyvacancies.com/assistantassociatefull-professor-of-computer-science,i43041.html.

Notice to authors