OSCAR 2026: The Fifth Workshop on Open-Source Computer Architecture Research

Luca Carloni

Department of Computer Science

Columbia University in the City of New York

A workshop co-located with ISCA 2026 --- Raleigh, NC --- June 28, 2026

OSCAR 2026

- OSCAR 2026 is the fifth edition of a workshop aimed at fostering the community of researchers who are interested in developing and sharing open-source hardware and software for the design of next-generation computer architectures
- Organizers:
  - Pradip Bose (IBM)
  - Luca Carloni (Columbia University), chair
  - Margaret Martonosi (Princeton University)
  - Sophia Shao (UC Berkeley)
  - Caroline Trippel (Stanford University)

Why OSCAR?

- The Premise
  - We are in the age of heterogeneous computing. This heterogeneity brings new challenges to hardware designers as well as software programmers
  - Addressing many of these challenges requires collaborative and open-source research
- The Rationale
  - Many innovations are best evaluated in the context of complete system-level implementations, which go beyond traditional simulation methods
  - Most individual research groups do not have the resources to realize such implementations

Why OSCAR?

- The Goal
  - Bring together a community of researchers from academia, industry and government labs who are interested in open-source computer architectures
  - The recent past has seen significant progress in this direction, including contributing open-source hardware components, software tools, as well as integration platforms to simplify the realization of system prototypes with FPGA or ASIC technologies. The number of developers and users of these open-source artifacts has increased substantially
  - It is time to provide a venue that promotes the growth of this community and fosters its efforts!

Today's Program - 1

8:00-8:20: **Opening Remarks** *(Luca Carloni)*

8:20-9:00: **Session I – Compilers and Runtime** *(Session Chair: Luca Carloni)*

1. Agustin Nahuel Coppari Hollmann, Kris Dong, Dima Nikoforov, Ashvin Verma and Sophia Yakun Shao
   Merlin: MLIR-Based Compiler and Runtime Generation for Hardware Integration
2. Shinnung Jeong, Chihyo Ahn, Sai Hemanth Reddy Bheemreddy,Liam Cooper, Krishil Gandhi, Chulhyung Park, Vincent Pham,Huanzhi Pu, Saurabh Singh, Mitul Tandon, Jisheng Zhao and Hyesoon Kim
   Toward End-to-End ML Support on Vortex

9:00-10:00: **Session I – Security** *(Session Chair: Luca Carloni)*

1. Kelly Farran, Derek Hansen, Michael Shires and Kevin Skadron
   Modeling Error Correction for PIM Architectures in PIMeval-PIMbench
2. Gaurav Kuwar, Alhad Daftardar, Jianqiao Mo and Brandon Reagen
   Tessera: An Open-Source Hardware Framework for Cryptographic Kernels
3. Dillibabu Shanmugam and Patrick Schaumont
   Open-Source Reference for Reproducible Side-Channel Assessment of Post-Quantum Standards

10:00-10:30: **Coffee Break**

Today's Program - 2

10:30-12:00: **Session III – Physical Design** *(Session Chair: Sophia Shao)*

1. Brayden Louie, Thinh Nguyen, Matt Liberty and Austin Rovinski
   ORRAM: An OpenROAD-Integrated RAM Generator Using Standard Cells
2. Arya Reais-Parsi and John Wawrzynek
   BFG: An Open-Source Full-Custom Silicon Compiler for High-Performance FPGA Fabrics
3. Rahul Kumar, Rohan Kumar, Ethan Wu, Allison Husain, Connor Lu, Viansa Schmulbach and Borivoje Nikolić
   SRAM22: An Open-Source SRAM Compiler for Skywater 130nm
4. Andre Green, Desvaun Drummond, Brandon Kim, Juhyun Do, Lawrence Rhee, Isabelle Hsu, Henry Wilson, Colin O'Brien, Sagar Karandikar and Borivoje Nikolic
   Beyond Make: DAG-Driven Physical Design and Intelligent Dependencies with SledgeHammer

12:00-13:30: **Lunch**

Today's Program - 3

13:30-14:30: **Session IV – GPU I** *(Session Chair: Caroline Trippel)*

1. Hansung Kim, Ruohan Yan, Shashank Anand, Joshua You, Man Shi, Christopher Fletcher and Yakun Sophia Shao
   Muon: An ASIC-optimized SIMT Core Designed for Register-heavy GPU Kernels
2. Giorgos Triantafyllou, Panagiotis-Eleftherios Eleftherakis, Konstantinos Iliakis, Blaise Tine and Sotirios Xydis
   RTL Design and Analysis of Sparse Tensor Cores for Unstructured Sparsity in RISC-V GPUs
3. Yanggon Kim, Xinle Song and Blaise Tine
   Thread-Scalable 2:4 Structured Sparsity on an Open RISC-V GPU

14:30-15:10: **Session V – Architecture Optimization** *(Session Chair: Caroline Trippel)*

1. Phaedra Sophia Curlin and Tamara Silbergleit Lehman
   Open-Source Cache Architectures in Chipyard
2. Victor Isachi, Angelo Garofalo, Francesco Conti and Davide Rossi
   Throughput and Area Impact of Synchronization in an Open-Source Massively-Parallel Manycore Platform

Today's Program - 4

15:10-15:30: **Poster Session Lightning Talks** *(Session Chair: Luca Carloni)*

15:30-16:30: **Poster Session with Coffee Break**

1. Xinle Song and Blaise Tine
   CycleFence: Precise Cycle-Level Profiling for RISC-V GPUs
2. Maria Zerva, Negin Mahani and Osman Unsal
   An RTL Fault Injection Framework & ECC Support for Reliability Analysis of GPU Register Files
3. Yimin Gao, Liangtao Dai, Junting Huo and Mircea Stan
   FlexNPU: A Compiler-Integrated FPGA Platform for Plug-and-Play AI Accelerator Research
4. Chengxuan Wang and Blaise Tine
   Supporting Async Barrier and TMA for Vortex GPGPU
5. Kaushik Chandana, Jay Imperatori, Tanmay Shukla, Justin Zhou, Stefan Abi-Karam and Callie Hao
   HLSFactory-Agent: Large-Scale Agentic HLS Dataset Construction from Academic and Open-Source Projects
6. Jude Haris, Rappy Saha and José Cano
   Hardware Profiling for SECDA on Edge Platforms

Today's Program - 5

7. Randy Liu and Blaise Tine
   Hardware Acceleration of GPU Kernel Scheduling on RISC-V GPU
8. William Bradford, Farzana Ahmed Siddique, Nebil Ozer, Kevin Skadron and Sandhya Dwarkadas
   Build-A-PIM: A Modular Framework for Exploring the "What-if's" of Processing in Memory
9. Nikhil Rout and Blaise Tine
   RoadBlock: Rethinking Tensor Core Microarchitecture for Emerging Microscaling Format Support on GPUs
10. Sahil Gandhi, Anirudh Jain, Jiyong Yu and Vignyan Kothinti
    RPM: RISC-V Performance Model
11. Ubaid Bakhtiar and Bahar Asgari
    Towards Transparent and Reproducible Accelerator Research
12. Shwet Chitnis, Jiayi Wang, Fergus Xu, Ayush Kulkarni, Rampranav Navendran, Juwon Jun, Jingqun Zhang and Ang Li
    DORA: Open-Source Infrastructure for Prototyping Reconfigurable Fabrics
13. Kuan Fu Chen and Blaise Tine
    PRISM: Accelerating Ray Tracing on RISC-V GPU
14. Injae Shin and Blaise Tine
    gem5-Vortex: Heterogeneous Cross-ISA Integration of a RISC-V GPGPU in gem5

Today's Program - 6

16:30-17:30: **Session VI – GPU II** *(Session Chair: Pradip Bose)*

1. Xilai Dai and Mohamed Abdelfattah
   Toward Sustainable GPU Modeling: Autonomous Micro-benchmarking with LLM Agents
2. Jianming Tong, Devansh Jain, Niansong Zhang, Charith Mendis, Zhiru Zhang and Tushar Krishna
   RAIC: Full-stack Reconfigurable AI Computing Generation
3. Michael Lippe, Luca Carloni, Biruk Seyoum, Gabriele Tombesi and Joseph Zuckerman
   An Integration of the Open-Source Vortex GPU into the ESP Platform

17:30-18:30: **Session VI – Platforms** *(Session Chair: Luca Carloni)*

1. Christian Pilato, Luis Miguel Pinho, Rene Scholz, Frank K. Gürkaynak, Luca Benini and Stefan Wallentowitz
   Open Design Flows for Open Hardware Research: The ODE4EC Perspective
2. Ankur Limaye, Nicolas Bohm Agostini, David Kong, Nrusinga Charan Gantayat, Gianmarco Accordi, Max Ramstad, Lakshmi Varshika Mirtinti, Vito Giovanni Castellana, Joseph Manzano, Jeff J. Zhang, Gage Hills, Fabrizio Ferrandi and Antonino Tumeo
   From Python to Silicon: First Tapeouts Produced by an End-to-End Open-Source Hardware Compiler
3. Pradip Bose, Luca Carloni, Kenneth Shepard, David Brooks, Gu-Yeon Wei, Vijay Janapa Reddi, Sarita Adve, Vikram Adve and Sasa Misailovic
   Early Stage SoC Architectural Concept Definition: Open-Source Toolset from the EPOCHS Project

18:30: **Concluding Remarks** *(Session Chair: Luca Carloni)*

OSCAR Program: Historical Statistics

OSCAR Edition	2022	2023	2024	2025	2026
Submissions	22	25	19	25	38
Accepted Talks	12	17	13	17	20
Accepted Posters	7	6	4	8	14

OSCAR Webpage Archive

Open-Source Computer Architecture Research (OSCAR)

OSCAR Workshop - Archive

- 2025 - Tokyo, Japan
- 2024 - Buenos Aires, Argentina
- 2023 - Orlando, FL
- 2022 - New York, NY

Why OSCAR after Five Years?

- Open-Source Hardware is even more relevant in the Age of Agentic AI
  - Shareable artifacts help everyone, smaller and larger teams, to complete more compelling experiments
  - Agentic AI for Computer Architecture and CAD need OSH for model development
  - OSH supports the design of new courses and curricula to develop the next generation of computer engineers (who think at the system level…)
  - The combination of OSH and AI enables more reproducible research evaluation
  - Sharing builds collaborations across academia, government labs, and industry
- Let's talk about this together today and keep build the OSCAR community!

Davide Giri (1990-2021)