

Hot Chips 2025 Program Committee

Thank you
to the
Program
Committee!

Program	Co-Chairs	Ian Bratt	ARM
		Nhon Quach	Samsung
Members		Ralph Wittig	AMD
		Greg Papadopoulos	NEA
		Pradeep Dubey	Intel
		Alan Jay Smith	UC Berkeley
		Gabriel Southern	Qualcomm
		Ronny Krashinsky	NVIDIA
		Borivoje Nikolic	UC Berkeley
		Yasuo Ishii	Tenstorrent
		Yuan Xie	Hong Kong University
		Jae W. Lee	Seoul National University
		Lavanya Subramanian	OpenAI
		Cliff Young	Google
		Sherry Xu	Microsoft
		Mahesh Maddury	Meta
		Fredrik Kjolstad	Stanford

Keynotes

Monday

2:15PM-
3:15PM

Keynote #1

Chair: Cliff Young

Predictions for the Next Phase of AI

Noam Shazeer, GDM VP,
Engineering, Google

Tuesday

1:45PM-
2:45PM

Keynote #2

Chair: Yasuo Ishii

**Up and Running with Rapidus: How Japan and Cutting-Edge
Technologies are Transforming Semiconductor Manufacturing**

Dr. Atsuyoshi Koike, Rapidus

Program (Monday)

Time	Session	Session Chair	Submission	Title	Speaker
9:15	Opening Remarks	Ian/Nhon/JW/Larry			
9:30	CPU	Gabriel Southern	46	Cuzco: A High-Performance RISC-V RVA23 Compatible CPU IP	Ty Garibay and Shashank Nemawarkar
10:00	CPU		22	PEZY-SC4s: The Fourth Generation MIMD Many-core Processor with High Energy Efficiency and Flexibility for AI	Naoya Hatta
10:30	CPU		23	IBM's Next Generation Power Microprocessor	William Starke
11:00	Break				
11:30	CPU		10	Introducing the Next Generation Intel® Xeon® Processor with Efficiency Cores	Don Soltis
12:00	Security	Greg Papadopoulos	68	Presto: A RISC-V-Compatible SoC for Unified Multi-Scheme FHE Acceleration over Module Lattice	Luchang Lei and Hongyang Jia
12:30	Security		80	Azure Secure Hardware Architecture: Establishing a Robust Security Foundation for Cloud Workloads	Bryan Kelly
13:00	Lunch				
14:15	Keynote 1	Cliff Young		Predictions for the Next Phase of AI	Noam Shazeer
15:15	Graphics	Lavanya Subraman	8	AMD RDNA 4 and Radeon RX 9000 Series GPU	Andy Pomianowski and Laks Pappu
15:45	Graphics		33	RTX 5090: Designed for the Age of Neural Rendering	Marc Blackstein
16:15	Graphics		91	Specialized SoC enabling low-power 'World Lock Rendering' in Augmented and Mixed Reality Devices	Ohad Meitav and Jay Tsao
16:45	Break				
17:15	Networking	Sherry Xu	19	Intel Mount Morgan Infrastructure Processing Unit (IPU)	Patrick Fleming & <TBD>
17:45	Networking		26	AMD Pensando™ Pollara 400 AI NIC Architecture and Application	Kevin Chu
18:15	Networking		69	NVIDIA ConnectX-8 SuperNIC: A Programmable RoCE Architecture for AI Data Centers	Idan Burstein
18:45	Networking		59	Tomahawk Ultra - Ultra Low Latency, High Bandwidth Ethernet Switch chip for HPC and AI/ML applications	Mohan Kalkunte & Asad Khamisy
19:15	Reception				

Program (Tuesday)

8:30	Optical	Borivoje Nikolic	76	Celestial AI Photonic Fabric Module (PF Module) - The world's first SoC with in-die Optical IO	Phil Winterbottom
9:00	Optical		83	A UCIe Optical I/O Retimer Chiplet for AI Scale-up Fabrics	Vladimir Stojanovic
9:30	Optical		57	Passage M1000: 3D photonic interposer for AI	Darius Bunandar
10:00	Optical		52	Co-Packaged Silicon Photonics Switches for Gigawatt AI Factories	Gilad Shainer
10:30	Break				
11:00	Power / Method	Jae W. Lee	62	ECAM Enabled Advanced Thermal Management Solutions for the AI Data Center	Michael Matthews
11:30	Power / Methodology		37	Everactive Self-Powered SoC with Energy Harvesting, Wakeup Receiver, and Energy-Aware Subsystem	Ben Calhoun
12:00	Power / Methodology		85	Taping Out Three Class Chips per Semester in Intel 16 Technology	Lucy Revina
12:30	Lunch				
13:40	TCMM			TCMM Awards Presentation	
13:45	Keynote 2	Yasuo Ishii		Up and Running with Rapidus: How Japan and Cutting-Edge Technologies are Transforming Semiconductor	Dr. Atsuyoshi Koike
14:45	Machine Learning	Ronny Krashinsky	40	Memory: (Almost) the Only Thing That Matters	Mark Kuemerle
15:15	Machine Learning 1		24	Corsair - An In-memory Computing Chiplet Architecture for Inference-time Compute Acceleration	Sudeep Bhoja
15:45	Machine Learning 1		67	UB-Mesh: Huawei's Next-Gen AI SuperComputer with A Unified-Bus Interconnect and nD-FullMesh Architecture	Liao Heng
16:15	break				
16:45	Machine Learning	Pradeep Dubey	21	NVIDIA's GB10 SoC: AI Supercomputer On Your Desk	Andi Skende
17:15	Machine Learning 2		6	4th Gen AMD CDNA™ Generative AI Architecture Powering AMD Instinct™ MI350 Series GPUs and Platform	Michael Floyd and Michael Steffen
17:45	Machine Learning 2		61	Ironwood : Delivering best in class perf, perf/TCO and perf/Watt for reasoning model training and serving	Norm Jouppi and Sridhar Lakshmanamurthy

Posters

Best poster (TCMM) award to
be handed out before Tuesday
Keynote

Title	Authors & Affiliation
HyperAccel Adelia: A 4nm LLM Processor for Efficient Generative AI Inference	Seungjae Moon; Hyper Accel
Basilisk: A 34 mm ² End-to-End Open-Source 64-bit Linux-Capable RISC-V SoC in 130nm BiCMOS	Philippe Sauter; ETH Zurich
Bit-Separable Transformer Accelerator Leveraging Output Activation Sparsity for Efficient DRAM Access	Seunghyun Park; Kyungpook National University
Multi-modal Few-step Diffusion Model Accelerator with Mixed-Precision and Reordered Group-Quantization for On-device Generative AI	Sangjin Kim; KAIST
An Energy-Efficient Spatial Computing SoC for Real-time Interactable-Rendering and Modeling with Surface-aware 3D Gaussian Splatting	Seokchan Song; KAIST
BROCA: A Low-power and Low-latency Conversational Agent RISC-V System-on-Chip for Voice-interactive Mobile Devices	Wooyoung Jo; KAIST
MEGA.mini: A NPU with Novel Heterogeneous AI Processing Architecture Balancing Efficiency, Performance, and Intelligence for the Era of Generative AI	Donghyeon Han; Chung-Ang University
High Density Si-IPD Technologies as enabler for High-Performance and Low-Power consumption Processor Chips	Mohamed Mehdi Jatlaoui; Murata
Clo-HDnn: Continual On-Device Learning Accelerator with Hyperdimensional Computing via Progressive Search	Chang Eun Song; UC San Diego
A 4.69mW LLM Processor with Binary/Ternary Weights for Billion-Parameter Llama Model	Sangyeob Kim; Yonsei University
KLIMA: Low-latency mixed-signal In-Memory Computing accelerator for solving arbitrary-order Boolean Satisfiability	Tinish Bhattacharya; UC Santa Barbara

Thanks!