

Welcome to Hot Chips 2023 Program

Natalia Vassilieva (Cerebras)

Heiner Litz (UC Santa Cruz)

Hot Chips 2023 Program Committee

Ralph Wittig, Xilinx/AMD

Pradeep K Dubey, Intel

Priyanka Raina, Stanford

Jae W. Lee, SNU

Forest Baskett, Nea

Mark D. Hill, Microsoft

Alan Smith, UC Berkeley

Ashkan Seyedi, Nvidia

Lingjie Xu

Debjit Das Sarma, Tesla

Yasuo Ishii, Tenstorrent

Tom St. John, OctoML

Grant Ayers, Google

Dave Ditzel, Esperanto

Tutorials

ML Inference

- Overview, Micah Villmow, NVIDIA
- Quantization Methods for Efficient ML Inference, Amir Gholami, UC Berkeley
- ML Inference at the Edge, Felix Baum, Qualcomm
- PyTorch 2.0, Elias Ellison, Meta
- Hardware Requirements for Exploiting Sparsity in ML Inference, Zhibin Xiao, Moffett AI

Chiptlets/UCI

- Overview, Debendra Das Sharma, Intel & Nathan Kalyansundharam, AMD
- UCle Protocol, Swadesh Choudhary, Intel & Marvin Denman, NVIDIA
- Electrical, Form Factor and Compliance, Anwar Kashem, AMD & Gerald Pasdast, Intel
- Software, Manageability & Security, Jérôme Glisse, Google & Sridhar Muthrasanallur, Intel

Keynotes



Exciting Directions for ML Models and the
Implications for Computing Hardware

Jeff Dean, Google, Chief Scientist
Amin Vahdat, Google, Fellow & VP

Monday, 9:00 am



Hardware for Deep Learning

Bill Dally, NVIDIA, Chief Scientist

Tuesday, 9:00 am

Program Statistics and Sessions

- 90 Submissions
- 24 Accepted oral presentations (8 Sessions)
- Best submissions will be invited to IEEE MICRO
- 10 Accepted posters

Monday

- Processing in Memory
- CPU 1
- Platforms
- CPU 2

Tuesday

- ML-Training
- Interconnects
- ML-Inference
- FPGA & Cooling

Posters

- Selected 10 Outstanding Posters
- Poster Chair: Priyanka Raina
- Posters located at coffee break area
- Posters will be up for the entire conference
- Poster Lightning talks at 8:30 Tuesday (before Keynote)

	Sunday, August 26	Monday, August 27	Tuesday, August 28
7:45am	Registration & Breakfast	Registration & Breakfast	Registration & Breakfast
10:00am	8:30am - 10:30am Tutorial 1: ML Inference	8:45am: Opening Remarks	8:30am: Poster Lightning Talks
		9:00am: Keynote 1 Jeff Dean & Amin Vahdat, Google	9:00am: Keynote 2 Bill Dally, NVIDIA
		10:00am - 11:00am Processing in Memory	10:00am - 11:00am ML Training
12:00pm	Coffee break & Posters	Sachs Memorial & Coffee break	Coffee break & Posters
	11:00am - 12:15pm Tutorial 1: ML Inference	11:30am - 1:00pm CPU 1	11:30am - 1:00pm Interconnects
	Lunch	Lunch	Lunch
2:00pm	1:30pm - 3:30pm Tutorial 2: Chiplets/UCI	2:15pm - 4:15pm Platforms	2:15pm - 4:15pm ML Inference
4:00pm	Coffee break & Posters	Coffee break & Posters	Coffee break & Posters
6:00pm	4:00pm - 5:30pm Tutorial 2: Chiplets/UCI	4:45pm - 6:15pm CPU 2	4:45pm - 6:15pm FPGAs & Cooling
	Reception	Reception	Closing Remarks