

Ashutosh Pattnaik

Email: ashutosh13@gmail.com

Homepage: <https://ashutoshpattnaik.github.io>

RESEARCH INTERESTS

GPU Architectures, Computer Graphics, Near-Data Computing, Machine Learning Architectures

EDUCATION

Doctor of Philosophy in Computer Science and Engineering Fall 2013 - Fall 2019
Pennsylvania State University, University Park, PA

Bachelor of Technology (Honors) in Electronics and Instrumentation Engineering Fall 2009 - Spring 2013
National Institute of Technology, Rourkela, India

WORK EXPERIENCE

Senior GPU Architect, Arm **September 2019 - Present**
Austin, TX

- Working on the architecture design and optimization of the Arm Mali GPU to improve performance, power and area.

Research Assistant, Pennsylvania State University Fall 2013 - Fall 2019
University Park, PA

- Understanding research issues and opportunities involved in bringing near-data computing paradigm to GPUs and CPU-GPU systems and optimizing the compute placement for improved performance and energy efficiency.
- Improving the GPU datapath for improved execution of irregular applications and machine learning applications.

Co-Op Engineer, AMD Research May 2016 - August 2016
Sunnyvale, CA

- Researched on efficient interconnect topologies and data placement techniques for a 3D-stacked processing-in-memory enabled, multi-chip module based architecture.

Co-Op Engineer, AMD Research June 2015 - September 2015
Austin, TX

- Developed methodology and micro-benchmarks to perform detailed characterization of the energy usage of different ISA instructions and data movement in AMD GPUs.

PUBLICATIONS

[ISCA 2020] Sonali Singh, Anup Sarma, Nicholas Jao, **Ashutosh Pattnaik**, Sen Lu, Kezhou Yang, Abhronil Sengupta, Vijaykrishnan Narayanan, Chita Das, "**NEBULA: A Neuromorphic Spin-Based Ultra-Low Power Architecture for SNNs and ANNs**", In the Proceedings of the 47th International Symposium on Computer Architecture, Worldwide Event, June 2020.

[MICRO 2019] Haibo Zhang, Shulin Zhao, **Ashutosh Pattnaik**, Mahmut T. Kandemir, Anand Sivasubramaniam and Chita R. Das, "**Distilling the Essence of Raw Video to Reduce Memory Usage and Energy at Edge Devices**", In the Proceedings of the 52nd International Symposium on Microarchitecture, Columbus, Ohio, October 2019. *Acceptance Rate $\approx 23\%$*

[ISCA 2019] **Ashutosh Pattnaik**, Xulong Tang, Onur Kayiran, Adwait Jog, Asit Mishra, Mahmut T. Kandemir, Anand Sivasubramaniam and Chita R. Das, "**Opportunistic Computing in GPU Architecture**", In the Proceedings of the 46th International Symposium on Computer Architecture, Phoenix, Arizona, June 2019. *Acceptance Rate $\approx 17\%$*

[MEMSYS 2019] Anup Sarma, Huaipan Jiang, **Ashutosh Pattnaik**, Jagadish Kotra, Mahmut T. Kandemir, Chita R. Das, "**CASH: Improving DRAM Energy Efficiency in CPU-based Inference**", In the Proceedings of the 5th International Symposium on Memory Systems, Washington D.C., October 2019.

[SIGMETRICS 2019] Xulong Tang, [Ashutosh Pattnaik](#), Onur Kayiran, Adwait Jog, Mahmut T. Kandemir and Chita R. Das, "Quantifying Data Locality in Dynamic Parallelism in GPUs", In the Proceedings of the ACM on Measurement and Analysis of Computing Systems, Phoenix, Arizona, June 2019. *Acceptance Rate* $\approx 8\%$

[HPCA 2017] Xulong Tang, [Ashutosh Pattnaik](#), Huaipan Jiang, Onur Kayiran, Adwait Jog, Sreepathi Pai, Mohamed Ibrahim, Mahmut T. Kandemir and Chita R. Das, "Controlled Kernel Launch for Dynamic Parallelism in GPUs" In the Proceedings of the 23rd International Symposium on High Performance Computer Architecture, Austin, Texas, February 2017. *Acceptance Rate* $\approx 22\%$

[PACT 2016] [Ashutosh Pattnaik](#), Xulong Tang, Adwait Jog, Onur Kayiran, Asit Mishra, Mahmut Kandemir, Onur Mutlu and Chita Das, "Scheduling Techniques for GPU Architectures with Processing-In-Memory Capabilities" In the Proceedings of the 25th Parallel Architecture and Compilation Techniques, Haifa, Israel, September 2016. *Acceptance Rate* $\approx 22.3\%$

[PACT 2016] Onur Kayiran, Adwait Jog, [Ashutosh Pattnaik](#), Rachata Ausavarungnirun, Xulong Tang, Mahmut Kandemir, Gabriel Loh, Onur Mutlu and Chita Das, " μ C-States: Fine-grained GPU Datapath Power Management" In the Proceedings of the 25th Parallel Architecture and Compilation Techniques, Haifa, Israel, September 2016. *Acceptance Rate* $\approx 22.3\%$

[IISWC 2016] Vignesh Adhinarayanan, Indrani Paul, Joseph Greathouse, Wei N. Huang, [Ashutosh Pattnaik](#) and Wuchun Feng, "Measuring and Modeling On-Chip Interconnect Power on Real Hardware", In the Proceedings of IEEE International Symposium on Workload Characterization, Providence, Rhode Island, 2016. (**Best Paper Award**) . *Acceptance Rate* $\approx 30.4\%$

[SIGMETRICS 2016] Adwait Jog, Onur Kayiran, [Ashutosh Pattnaik](#), Mahmut Kandemir, Onur Mutlu, Ravi Iyer and Chita Das, "Exploiting Core-Criticality for Enhanced Performance in GPUs", In the Proceedings of the 42nd ACM International Conference on Measurement and Modeling of Computer Systems, Antibes Juan-les-Pins, France, June 2016. *Acceptance Rate* $\approx 13.4\%$

[MEMSYS 2015] Adwait Jog, Onur Kayiran, Tuba Kesten, [Ashutosh Pattnaik](#), Evgeny Bolotin, Nilardish Chatterjee, Steve Keckler, Mahmut Kandemir and Chita Das, "Anatomy of GPU Memory System for Multi-Application Execution", In the Proceedings of the 1st International Symposium on Memory Systems, Washington D.C., October 2015.

[ICCCS 2012] [Ashutosh Pattnaik](#), Sharad Agarwal, Subhasis Chand, "A New and Efficient Method for Removal of High Density Salt and Pepper Noise Through Cascade Decision based Filtering Algorithm", In the Proceedings of the 2nd International Conference on Communication, Computing & Security, India, 2012.

TEACHING EXPERIENCE

Teaching Assistant, Penn State Spring 2014
CMPEN 431, Introduction to Computer Architecture

Teaching Assistant, Penn State Fall 2013
CMPEN 270, Digital Design: Theory and Practice

Guest Lecturer, Penn State
CSE 597: Advances and Applications in Deep Learning (Spring 2017)
CSE 532: Multiprocessor Architecture (Spring 2015)
CMPEN 431: Introduction to Computer Architecture (Fall 2017, Spring 2018)
CMPEN 331: Computer Organization And Design (Spring 2015)

TALKS

- Be(-A)ware of Data Movement: Optimizing Throughput Processors for Efficient Computations
Qualcomm, San Diego, CA, USA, July 2019

- Opportunistic Computing in GPU Architecture
ISCA 2019, Phoenix, AZ, USA, June 2019
- Quantifying Data Locality in Dynamic Parallelism in GPUs
SIGMETRICS 2019, Phoenix, AZ, USA, June 2019
- Be(-A)ware of Data Movement: Optimizing Throughput Processors for Efficient Computations
Advanced Micro Devices (AMD) Research, Austin, TX, USA, March 2019
- Scheduling Techniques for GPU Architectures with Processing-In-Memory Capabilities
PACT 2016, Haifa, Israel, September 2016
- μ C-States: Fine-grained GPU Datapath Power Management
PACT 2016, Haifa, Israel, September 2016
- Exploiting Core-Criticality for Enhanced Performance in GPUs
SIGMETRICS 2016, Antibes Juan-les-Pins, France, June 2016
- A New and Efficient Method for Removal of High Density Salt and Pepper Noise Through Cascade Decision based Filtering Algorithm
ICCCS 2012, India, October 2012

HONORS AND AWARDS

Best Paper Award:

- Measuring and Modeling On-Chip Interconnect Power on Real Hardware, IISWC 2016

Student Travel Grant:

- 2019: IEEE Travel Grant for ISCA, NSF Travel Grant for SIGMETRICS
- 2017: IEEE Travel Grant for HPCA, ACM SIGARCH Travel Grant for ISCA
- 2016: ACM SIGMETRICS Travel Grant, NSF Travel Grant for PACT
- 2015: ACM SIGARCH Travel Grant for ISCA

SERVICE AND MEMBERSHIPS

Session Chair: GPUs, IEEE International Symposium on High-Performance Computer Architecture (HPCA), San Diego, CA, Feb, 2020

External Review Committee Member, IEEE International Symposium on High-Performance Computer Architecture (HPCA), San Diego, CA, Feb, 2020

General Co-Chair, Workshop on General Purpose Processing Using GPU (GPGPU), San Diego, CA, Feb, 2020

Submission/Web Chair, Workshop on General Purpose Processing Using GPU (GPGPU), Providence, RI, April 2019

Submission Chair, International Conference on Supercomputing (ICS), Turkey, June 2016

Reviewer:

- Transactions on Cloud Computing, IEEE
- Transactions on Parallel and Distributed Systems, IEEE
- IEEE Transactions on Architecture and Code Optimization (TACO)
- IEEE Transactions on Computers (TC)
- IEEE Micro
- Microprocessors and Microsystems: Embedded Hardware Design, Elsevier
- ETRI Journal, Wiley

Memberships:

- ACM Professional, IEEE Young Professional, SIGARCH, IEEE Computer Society

REFERENCES

References are available on request.