

GPU synchronization (September 12):

Ahmed ElTantawy and Tor M. Aamodt
MIMD synchronization on SIMT architectures
MICRO, 2016
<https://ece.ubc.ca/~aamodt/papers/eltantawy.micro2016.pdf>

Jiwei (presenter), [Seyed](#), [John](#)

Daniel Lustig and Margaret Martonosi
Reducing GPU offload latency via fine-grained CPU-GPU synchronization
HPCA, 2013
<http://mrmgroup.cs.princeton.edu/papers/dlustigHPCA13.pdf>

Debashis (presenter), [Zaeem](#), [Injung](#)

Wu-chun Feng and Shucaï Xiao
To GPU Synchronize or Not GPU Synchronize?
International symposium on circuits and systems - 2010
<https://research.cs.vt.edu/synergy/pubs/papers/feng-iscas2010-gpusync.pdf>

Shucaï Xiao and Wu-chun Feng
Inter-Block GPU Communication via Fast Barrier Synchronization
IPDPS 2010
http://eprints.cs.vt.edu/archive/00001087/01/TR_GPU_synchronization.pdf

Mohamed (presenter), [Junda](#), [Kenrick](#)

GPU memory and cache management (September 19):

Mohammad Dashti and Alexandra Fedorova
Analyzing memory management methods on integrated CPU-GPU systems
Proceedings of the 2017 ACM SIGPLAN International Symposium on Memory Management
<https://www.ece.ubc.ca/~sasha/papers/ismm-2017.pdf>

Injung(presenter), [Debashis](#), [Zaeem](#)

Gunjae Koo, Yunho Oh, Won Woo Roand Murali Annavaram*
Access Pattern-Aware Cache Management for Improving Data Utilization in GPU
ISCA 2017
<http://dl.acm.org/citation.cfm?doid=3079856.3080239>

Seyed(presenter), [John](#), [Jiwei](#)

Xinxin Mei and Xiaowen Chu
Dissecting GPU Memory Hierarchy through Microbenchmarking
IEEE Transactions on Parallel and Distributed Systems · 2017
<https://arxiv.org/pdf/1509.02308.pdf>

Junda(presenter), [Kenrick](#), [Mohammad](#)

Fault recovery in HPC: (September 26)

Xiang Ni, Esteban Meneses, Nikhil Jain and Laxmikant V. Kalé
ACR: Automatic checkpoint/restart for soft and hard error protection
SC'13
http://charm.cs.uiuc.edu/users/xiang/SC_ACR.pdf

J.Elliott, K.Kharbas, D.Fiala, F.Mueller, K.Ferreira, C.Engelmann John(presenter), [Jiwei](#), [Seyed](#)
Combining Partial Redundancy and Checkpointing for HPC
ICDCS'12
<https://www.fiala.me/pubs/papers/icdcs12-redundancymodeling.pdf>

Shen Gao, Bingsheng He and Jianliang Xu
Real-Time In-Memory Checkpointing for Future Hybrid Memory Systems
SC'15
<http://www.comp.nus.edu.sg/~hebs/pub/DRAMCheckpoint-ICS15.pdf>

B.Fang, Q.Guan, N.Debardleben, K.Pattabiraman, M.Ripeanu Kenrick(presenter), [Junda](#), [Mohamed](#)
LetGo: A Lightweight Continuous Framework for HPC Applications under Failures.
International Symposium on High-Performance Parallel and Distributed computing(HPDC '17)
<https://doi.org/10.1145/3078597.3078609>

Gengbin Zheng, Xiang Ni and Laxmikant V. Kale Zaeem(presenter), [Debashis](#), [Injung](#)
A Scalable Double In-memory Checkpoint and Restart Scheme towards Exascale
Dependable systems and networks · 2012
<http://charm.cs.illinois.edu/newPapers/12-12/paper.pdf>

Fault detection in HPC: (October 3)

George Bosilca et al. Zaeem(presenter), [Debashis](#), [Injung](#)
Failure detection and propagation in HPC systems
SC'16
<https://hal.inria.fr/hal-01352109/file/sc16-hal.pdf>

Fiala, Mueller, Engelmann, Riesen, Ferreira and Brightwell Mohamed (presenter), [Junda](#), [Kenrick](#)
Detection and Correction of Silent Data Corruption for Large-Scale High-Performance Computing
SC'12
<http://moss.csc.ncsu.edu/~mueller/ftp/pub/mueller/papers/sc12.pdf>

Sheng Di and Franck Cappello Jiwei (presenter), Seyed, [John](#)
Adaptive Impact-Driven Detection of Silent Data Corruption for HPC Applications
IEEE Transactions on Parallel and Distributed Systems Volume: 27, Issue: 10, Oct. 2016
<http://www.mcs.anl.gov/papers/P5376-0715.pdf>

Amogh Katti, Giuseppe Di Fatta, Thomas Naughton and Christian Engelmann
Scalable and Fault Tolerant Failure Detection and Consensus
EuroMPI 15
<http://www.christian-engelmann.info/publications/katti15scalable.pdf>

Graph processing on GPU and shared memory (October 24 and Nov 7)

W. Zhong, J. Sun, H. Chen, J. Xiao, Z. Chen, C. Chang, and X. Shi

Optimizing Graph Processing on GPUs,

IEEE Transactions on Parallel and Distributed Systems (Volume: 28, Issue: 4, April 1 2017)

<https://www.aimlab.org/haochen/papers/tpds17-GPUgraph.pdf>

S. Julian and G. Blleloch.

Debashis (presenter), [Zaeem](#), [Injung](#)

Ligra: a lightweight graph processing framework for shared memory.

ACM Sigplan Notices. Vol. 48. No. 8. ACM, 2013.

<http://www.cs.cmu.edu/~jshun/ligra.pdf>

Y. Wang et al.

John (presenter), [Jiwei](#), [Seyed](#)

Gunrock: A high-performance graph processing library on the GPU.

ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming. ACM, 2016.

<https://arxiv.org/pdf/1501.05387v2.pdf>

Tao Zhang, Jingjie Zhang, Wei Shu, Min-You Wu and Xiaoyao Liang

Efficient graph computation on hybrid CPU and GPU systems

The Journal of Supercomputing, Volume 71 Issue 4, April 2015

<https://link.springer.com/article/10.1007/s11227-015-1378-z>

J. Malicevic, B. Lepers and W. Zwaenepoel

Kenrick (presenter), [Junda](#), [Mohammad](#)

Everything you wanted to know about multi-core graph processing but were afraid to ask

EPFL. USENIX ATC'17

<https://infoscience.epfl.ch/record/228854/files/atc17-final234.pdf>

Papers related to the Channel Attack project (October 31)

Zhen Hang Jiang, Yunsi Fei, David Kaeli

Seyed (presenter), [Jiwei](#), [John](#)

A Complete Key Recovery Timing Attack on a GPU

HPCA 2016

http://tescase.coe.neu.edu/uploaded_file/923be6c156a050543d1f92d8fce95f35a5bb7525.pdf

K. Khasawneh et al.

Junda (presenter), [Kenrick](#), [Mohammad](#)

RHMD: Evasion-Resilient Hardware Malware Detectors

Micro 2017

<http://www.cs.ucr.edu/~nael/pubs/micro17-rhmd.pdf>

J. Demme et Al.

On the Feasibility of Online Malware Detection with Performance Counters

ISCA 2013

http://www.cs.columbia.edu/~jdd/papers/isca13_malware.pdf

Z. He and R. Lee,

Injung (presenter), [Zaeem](#), [Debashis](#)

How secure is your cache against side-channel attacks?

MICRO-17

<https://dl.acm.org/citation.cfm?id=3124546>

Papers related to the GPU QoS project (November 14):

Tianhao Zheng et al. Debashis (presenter), [Zaeem](#), [Injung](#)
Towards High Performance Paged Memory for GPUs
HPCA 2016
https://www.cs.utexas.edu/~skeckler/pubs/HPCA_2016_Paged_Memory.pdf

B. Pichai, L. Hsu, and A. Bhattacharjee,
Architectural Support for Address Translation on GPUs: Designing Memory Management Units for
CPU/GPUs with Unified Address Spaces
ASPLOS 2014
<https://www.cs.rutgers.edu/~abhib/bpichai-asplos14.pdf>

Ali Bakhoda et al. Mohamed (presenter), [Junda](#), [Kenrick](#)
Analyzing CUDA Workloads Using a Detailed GPU Simulator
ISPASS 2009
https://john.cs.olemiss.edu/heroes/analyzingCUDAworkloadsUsingGPUsim_p.pdf

Rachata Ausavarungrun et al. Jiwei (presenter), [Seyed](#), [John](#)
Mosaic: A GPU Memory Manager with Application-Transparent Support for Multiple Page Sizes
MICRO 2017
<https://users.ece.cmu.edu/~rausavar/pubs/mosaic-micro17.pdf>

Cloud Processing (November 28):

S. Shastri and D. Irwina John (presenter), [Seyed](#), [Jiwei](#)
HotSpot: Automated Server Hopping in Cloud Spot Markets,
2017 ACM Symposium on Cloud Computing (SOCC 2017),
<http://www.ecs.umass.edu/~irwin/hotspot.pdf>

K. Tamrakar, A. Yazidi and H. Haugerud Kenrick (presenter), [Junda](#), [Mohamed](#)
Cost Efficient Batch Processing in Amazon Cloud with Deadline Awareness,
2017 IEEE 31st Int'l Conference on Advanced Information Networking and Applications,
<http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7921011>

S. Lo and Y. Cheng Injung (presenter), [Zaeem](#), [Debashis](#)
Improving the Performance of Fair Scheduler in Hadoop",
2017 Advanced in Science, technology and engineering systems journal (ASTES)
https://www.astesj.com/publications/ASTESJ_0203133.pdf

Miscellaneous: (December 5)

Jie Chen and Guru Venkataramani

Seyed(presenter), John, Jiwei

CC-Hunter: Uncovering Covert Timing Channels on Shared Processor Hardware

MICRO-2014

<https://pdfs.semanticscholar.org/21dd/f1f7ab7e2cd2ae07073bf3238ce46314bac9.pdf>

Jean-Pierre Lozi et al.

Zaeem(presenter), Injung, Debashis

The Linux Scheduler: a Decade of Wasted Cores

EuroSys 2016

<https://www.ece.ubc.ca/~sasha/papers/eurosys16-final29.pdf>

A. Milluzza and A. George

Junda (presenter), Kenrick, Mohamed

Exploration of TMR Fault Masking with Persistent Threads on Tegra GPU SoCs

Aerospace Conference, 2017

<http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7943882>

Distributed graph processing:

Ching, Avery, et al.

One trillion edges: Graph processing at facebook-scale.
Proceedings of the VLDB Endowment 8.12 (2015): 1804-1815.
<http://www.vldb.org/pvldb/vol8/p1804-ching.pdf>

Zhu, Xiaowei, et al.

Gemini: A Computation-Centric Distributed Graph Processing System.
OSDI. 2016.
<https://www.usenix.org/system/files/conference/osdi16/osdi16-zhu.pdf>

Malewicz, Grzegorz, et al.

Pregel: a system for large-scale graph processing.
Proceedings of the 2010 ACM SIGMOD International Conference on Management of data. ACM, 2010.
http://kowshik.github.io/JPregel/pregel_paper.pdf

Gonzalez, Joseph E., et al.

PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs.
OSDI. Vol. 12. No. 1. 2012.
<https://www.usenix.org/system/files/conference/osdi12/osdi12-final-167.pdf>

Accelerators:

Hadi Esmaeilzadeh, Adrian Sampson, Luis Ceze, and Doug Burger.
Neural Acceleration for General-Purpose Approximate Programs.
MICRO -2012
<http://dx.doi.org/10.1109/MICRO.2012.48>

Ham, Tae Jun, et al.

Graphicionado: A high-performance and energy-efficient accelerator for graph analytics.
MICRO, 2016.
http://mrmgroup.cs.princeton.edu/papers/taejun_micro16.pdf

T. Vijayaraghavan et al.

Design and Analysis of an APU for Exascale Computing
HPCA 2017
http://www.computermachines.org/joe/publications/pdfs/hpca2017_exascale_apu.pdf

Memory reliability:

Ruohuang Zheng and Michael C. Huang
Redundant Memory Array Architecture for Efficient Selective Protection
ISCA 2017
<http://www.ece.rochester.edu/~mihuang/PAPERS/isca17.pdf>

Xun Jian, Vilas Sridharan and Rakesh Kumar

Parity Helix: Efficient Protection for Single-Dimensional Faults in Multi-dimensional Memory Systems

HPCA 2016

http://rakeshk.crhc.illinois.edu/hpca_16_cam.pdf

Jungrae Kim, Michael Sullivan, Sangkug Lym and Mattan Erez

All-Inclusive ECC: Thorough End-to-End Protection for Reliable Computer Memory

ISCA 2016

<http://ieeexplore.ieee.org/document/7551427/>

Dong Wan Kim and Mattan Erez

RelaxFault Memory Repair

ISCA 2016

<http://ieeexplore.ieee.org/document/7551429/>