

# A Data-Intensive Computing Reading Group

University of Chicago, Statistics Department

October 4, 2015

## Purpose

As the importance of data intensive methods and applications grows, developing and implementing such methods is dependent on understanding the state of the art of data intensive computing. The goal of this reading group is to understand the historical and contemporary developments of data intensive computing so that it may inform the work we do in statistics, numerical methods, and machine learning.

## Reading Group Meetings

Meetings will be held weekly with two individuals presenting a paper per session. Presentations should be kept short (at most 20 minutes), provide sufficient background, and a summary of the work. All readings are mandatory for reading group members, and all meetings are mandatory (except for exams, conferences, illnesses, etc.) for all reading group members.

## Bibliography

Below is a working list of readings. This list is not set in stone and we can read and skip material as we see fit.

## Thermodynamics

- [1] Rolf Landauer. “Irreversibility and heat generation in the computing process”. In: *IBM journal of research and development* 5.3 (1961), pp. 183–191.
- [65] LV Zhirnov, Ralph Cavin, and Luca Gammaitoni. “Minimum energy of computing, fundamental considerations”. In: *ICTEnergyConcepts Towards Zero-Power Information and Communication Technology* 7 (2014).

## Paradigms

- [12] Jarek Nieplocha et al. “Advances, applications and performance of the global arrays shared memory programming toolkit”. In: *International Journal of High Performance Computing Applications* 20.2 (2006), pp. 203–231.
- [31] Michael G Burke et al. “Concurrent collections programming model”. In: *Encyclopedia of Parallel Computing*. Springer, 2011, pp. 364–371.
- [43] Jinsuk Chung et al. “Containment domains: A scalable, efficient and flexible resilience scheme for exascale systems”. In: *Scientific Programming* 21.3-4 (2013), pp. 197–212.

## Streaming Processing Systems

- [7] Daniel J Abadi et al. “Aurora: a new model and architecture for data stream management”. In: *The VLDB JournalThe International Journal on Very Large Data Bases* 12.2 (2003), pp. 120–139.
- [10] Daniel J Abadi et al. “The Design of the Borealis Stream Processing Engine.” In: *CIDR*. Vol. 5. 2005, pp. 277–289.
- [27] Leonardo Neumeyer et al. “S4: Distributed stream computing platform”. In: *Data Mining Workshops (ICDMW), 2010 IEEE International Conference on*. IEEE. 2010, pp. 170–177.
- [37] Gianpaolo Cugola and Alessandro Margara. “Processing flows of information: From data stream to complex event processing”. In: *ACM Computing Surveys (CSUR)* 44.3 (2012), p. 15.
- [45] Supun Kamburugamuve et al. *Survey of distributed stream processing for large stream sources*. Tech. rep. Technical report. 2013. Available at [http://grids.ucs.indiana.edu/ptliupages/publications/survey\\_stream\\_processing.pdf](http://grids.ucs.indiana.edu/ptliupages/publications/survey_stream_processing.pdf), 2013.

## Graph Processing Systems

- [13] Andrew Lumsdaine et al. “Challenges in parallel graph processing”. In: *Parallel Processing Letters* 17.01 (2007), pp. 5–20.
- [24] Grzegorz Malewicz et al. “Pregel: a system for large-scale graph processing”. In: *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*. ACM. 2010, pp. 135–146.
- [39] Joseph E Gonzalez et al. “PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs.” In: *OSDI*. Vol. 12. 1. 2012, p. 2.
- [40] Yucheng Low et al. “Distributed GraphLab: a framework for machine learning and data mining in the cloud”. In: *Proceedings of the VLDB Endowment* 5.8 (2012), pp. 716–727.
- [53] Reynold S Xin et al. “Graphx: A resilient distributed graph system on spark”. In: *First International Workshop on Graph Data Management Experiences and Systems*. ACM. 2013, p. 2.

- [58] Yucheng Low et al. “Graphlab: A new framework for parallel machine learning”. In: *arXiv preprint arXiv:1408.2041* (2014).

## Machine Learning

- [46] Tim Kraska et al. “MLbase: A Distributed Machine-learning System.” In: *CIDR*. 2013.
- [49] Evan R Sparks et al. “MLI: An API for distributed machine learning”. In: *Data Mining (ICDM), 2013 IEEE 13th International Conference on*. IEEE. 2013, pp. 1187–1192.

## Numerical Methods

- [2] Sivan Toledo. “A survey of out-of-core algorithms in numerical linear algebra”. In: *External Memory Algorithms and Visualization* 50 (1999), pp. 161–179.
- [4] Eran Rabani and Sivan Toledo. “Out-of-Core SVD and QR Decompositions.” In: *PPSC*. 2001.
- [5] Yen-Yu Chen, Qingqing Gan, and Torsten Suel. “I/O-efficient techniques for computing PageRank”. In: *Proceedings of the eleventh international conference on Information and knowledge management*. ACM. 2002, pp. 549–557.
- [11] Mario Rosario Guarracino, Francesca Perla, and Paolo Zanetti. “A parallel block Lanczos algorithm and its implementation for the evaluation of some eigenvalues of large sparse symmetric matrices on multicomputers”. In: *Int. J. Appl. Math. Comput. Sci* 16.2 (2006), pp. 241–249.
- [56] James Elliott, Mark Hoemmen, and Frank Mueller. “Resilience in numerical methods: A position on fault models and methodologies”. In: *arXiv preprint arXiv:1401.3013* (2014).

## Parallel Processing Engines

- [17] Jeffrey Dean and Sanjay Ghemawat. “MapReduce: simplified data processing on large clusters”. In: *Communications of the ACM* 51.1 (2008), pp. 107–113.
- [18] Ralf Lämmel. “Googles MapReduce programming modelRevisited”. In: *Science of computer programming* 70.1 (2008), pp. 1–30.
- [21] Daniel Warneke and Odej Kao. “Nephele: efficient parallel data processing in the cloud”. In: *Proceedings of the 2nd workshop on many-task computing on grids and supercomputers*. ACM. 2009, p. 8.
- [22] Dominic Battré et al. “Nephele/PACTs: a programming model and execution framework for web-scale analytical processing”. In: *Proceedings of the 1st ACM symposium on Cloud computing*. ACM. 2010, pp. 119–130.
- [23] Jeffrey Dean and Sanjay Ghemawat. “MapReduce: a flexible data processing tool”. In: *Communications of the ACM* 53.1 (2010), pp. 72–77.

- [26] Sergey Melnik et al. “Dremel: interactive analysis of web-scale datasets”. In: *Proceedings of the VLDB Endowment* 3.1-2 (2010), pp. 330–339.
- [30] Matei Zaharia et al. “Spark: cluster computing with working sets”. In: *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*. Vol. 10. 2010, p. 10.
- [32] Sergey Bykov et al. “Orleans: cloud computing for everyone”. In: *Proceedings of the 2nd ACM Symposium on Cloud Computing*. ACM. 2011, p. 16.
- [41] Justin M Wozniak et al. “Turbine: A distributed-memory dataflow engine for extreme-scale many-task applications”. In: *Proceedings of the 1st ACM SIGMOD Workshop on Scalable Workflow Execution Engines and Technologies*. ACM. 2012, p. 5.
- [52] Justin M Wozniak et al. “Swift/T: large-scale application composition via distributed-memory dataflow processing”. In: *Cluster, Cloud and Grid Computing (CCGrid), 2013 13th IEEE/ACM International Symposium on*. IEEE. 2013, pp. 95–102.
- [55] Timothy G Armstrong et al. “Compiler techniques for massively scalable implicit task parallelism”. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Press. 2014, pp. 299–310.
- [57] Scott J Krieder et al. “Design and evaluation of the gemitc framework for GPU-enabled many-task computing”. In: *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing*. ACM. 2014, pp. 153–164.

## Resource Management Systems

- [33] Ali Ghodsi et al. “Dominant Resource Fairness: Fair Allocation of Multiple Resource Types.” In: *NSDI*. Vol. 11. 2011, pp. 24–24.
- [34] Benjamin Hindman et al. “Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center.” In: *NSDI*. Vol. 11. 2011, pp. 22–22.
- [35] Arun Raman et al. “Sprint: speculative prefetching of remote data”. In: *ACM SIGPLAN Notices*. Vol. 46. 10. ACM. 2011, pp. 259–274.
- [36] Zhiming Shen et al. “Cloudscale: elastic resource scaling for multi-tenant cloud systems”. In: *Proceedings of the 2nd ACM Symposium on Cloud Computing*. ACM. 2011, p. 5.
- [48] Kay Ousterhout et al. “Sparrow: distributed, low latency scheduling”. In: *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. ACM. 2013, pp. 69–84.
- [50] Vinod Kumar Vavilapalli et al. “Apache hadoop yarn: Yet another resource negotiator”. In: *Proceedings of the 4th annual Symposium on Cloud Computing*. ACM. 2013, p. 5.
- [51] Ke Wang, Kevin Brandstatter, and Ioan Raicu. “SimMatrix: SIMulator for MAny-Task computing execution fabRIC at eXascale”. In: *Proceedings of the High Performance Computing Symposium*. Society for Computer Simulation International. 2013, p. 9.

- [59] Iman Sadooghi et al. “Achieving efficient distributed scheduling with message queues in the cloud for many-task computing and high-performance computing”. In: *Cluster, Cloud and Grid Computing (CCGrid), 2014 14th IEEE/ACM International Symposium on*. IEEE. 2014, pp. 404–413.
- [60] Ke Wang et al. “Next generation job management systems for extreme-scale ensemble computing”. In: *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing*. ACM. 2014, pp. 111–114.
- [61] Ke Wang et al. “Optimizing load balancing and data-locality with data-aware scheduling”. In: *Big Data (Big Data), 2014 IEEE International Conference on*. IEEE. 2014, pp. 119–128.

## Storage Systems

- [3] Robert B Ross, Rajeev Thakur, et al. “PVFS: A parallel file system for Linux clusters”. In: *Proceedings of the 4th annual Linux Showcase and Conference*. 2000, pp. 391–430.
- [6] Frank B Schmuck and Roger L Haskin. “GPFS: A Shared-Disk File System for Large Computing Clusters.” In: *FAST*. Vol. 2. 2002, p. 19.
- [8] S Donovan et al. “Lustre: Building a file system for 1000-node clusters”. In: *Proceedings of the Linux Symposium*. 2003.
- [9] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. “The Google file system”. In: *ACM SIGOPS operating systems review*. Vol. 37. 5. ACM. 2003, pp. 29–43.
- [14] Eduardo Pinheiro, Wolf-Dietrich Weber, and Luiz André Barroso. “Failure Trends in a Large Disk Drive Population.” In: *FAST*. Vol. 7. 2007, pp. 17–23.
- [15] SA Weil et al. “Ceph: a scalable, high-performance distributed file system”. In: *OSDI06 Proceedings of the 7th symposium on operating systems design and implementation, Berkeley, CA*. 2007.
- [16] Fay Chang et al. “Bigtable: A distributed storage system for structured data”. In: *ACM Transactions on Computer Systems (TOCS)* 26.2 (2008), p. 4.
- [19] Brent Welch et al. “Scalable Performance of the Panasas Parallel File System.” In: *FAST*. Vol. 8. 2008, pp. 1–17.
- [20] Samuel Lang et al. “I/O performance challenges at leadership scale”. In: *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*. ACM. 2009, p. 40.
- [25] Carlos Maltzahn et al. “Ceph as a scalable alternative to the Hadoop Distributed File System”. In: *login: The USENIX Magazine* 35 (2010), pp. 38–49.
- [28] Konstantin Shvachko et al. “The hadoop distributed file system”. In: *Mass Storage Systems and Technologies (MSST), 2010 IEEE 26th Symposium on*. IEEE. 2010, pp. 1–10.

- [29] Ashish Thusoo et al. “Hive-a petabyte scale data warehouse using hadoop”. In: *Data Engineering (ICDE), 2010 IEEE 26th International Conference on*. IEEE. 2010, pp. 996–1005.
- [38] Cliff Engle et al. “Shark: fast data analysis using coarse-grained distributed memory”. In: *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. ACM. 2012, pp. 689–692.
- [42] Matei Zaharia et al. “Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing”. In: *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*. USENIX Association. 2012, pp. 2–2.
- [44] James C Corbett et al. “Spanner: Googles globally distributed database”. In: *ACM Transactions on Computer Systems (TOCS)* 31.3 (2013), p. 8.
- [47] Tonglin Li et al. “ZHT: A light-weight reliable persistent dynamic scalable zero-hop distributed hash table”. In: *Parallel & Distributed Processing (IPDPS), 2013 IEEE 27th International Symposium on*. IEEE. 2013, pp. 775–787.
- [54] Dongfang Zhao et al. “Distributed data provenance for large-scale data-intensive computing”. In: *Cluster Computing (CLUSTER), 2013 IEEE International Conference on*. IEEE. 2013, pp. 1–8.
- [62] Dongfang Zhao, Kan Qiao, and Ioan Raicu. “Hycache+: Towards scalable high-performance caching middleware for parallel file systems”. In: *Cluster, Cloud and Grid Computing (CCGrid), 2014 14th IEEE/ACM International Symposium on*. IEEE. 2014, pp. 267–276.
- [63] Dongfang Zhao et al. “FusionFS: Toward supporting data-intensive scientific applications on extreme-scale high-performance computing systems”. In: *Big Data (Big Data), 2014 IEEE International Conference on*. IEEE. 2014, pp. 61–70.
- [64] Dongfang Zhao et al. “Virtual chunks: On supporting random accesses to scientific data in compressible storage systems”. In: *Big Data (Big Data), 2014 IEEE International Conference on*. IEEE. 2014, pp. 231–240.