



















- [6]. J. Yang, H. Cui, J. Wu, Y. Tang, and G. Hu, "Determinism Is Not Enough : Making Parallel Programs Reliable with Stable Multithreading," *Communication of the ACM 2014 (2013)*.
- [7]. M. Garland, M. Kudlur, and Y. Zheng, "Designing a unified programming model for heterogeneous machines," in *2012 International Conference for High Performance Computing, Networking, Storage and Analysis*, 2012, pp. 1–11.
- [8]. F. Ji and X. Ma, "Using Shared Memory to Accelerate MapReduce on Graphics Processing Units," in *2011 IEEE International Parallel & Distributed Processing Symposium*, 2011, pp. 805–816.
- [9]. W. Fang, B. He, Q. Luo, N. K. Govindaraju, and T. Wang, "Mars : Accelerating MapReduce with Graphics Processors," *IEEE Trans. Parallel Distrib. Syst.*, vol. 22, no. 4, pp. 608–620, 2011.
- [10]. J. a. Stuart and J. D. Owens, "Multi-GPU MapReduce on GPU Clusters," in *2011 IEEE International Parallel & Distributed Processing Symposium*, 2011, pp. 1068–1079.
- [11]. C. Ranger, R. Raghuraman, A. Penmetta, G. Bradski, and C. Kozyrakis, "Evaluating MapReduce for Multi-core and Multiprocessor Systems," in *2007 IEEE 13th International Symposium on High Performance Computer Architecture*, 2007, pp. 13–24.
- [12]. R. M. Yoo, A. Romano, and C. Kozyrakis, "Phoenix rebirth: Scalable MapReduce on a large-scale shared-memory system," in *2009 IEEE International Symposium on Workload Characterization (IISWC)*, 2009, pp. 198–207.
- [13]. J. Talbot, R. M. Yoo, and C. Kozyrakis, "Phoenix ++ : Modular MapReduce for Shared-Memory Systems," in *In Proceedings of the second international workshop on MapReduce and its applications (MapReduce '11)*, 2011, pp. 9–16.
- [14]. T. Liu and E. D. Berger, "D THREADS : Efficient Deterministic Multithreading," in *In Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles (SOSP '11)*, 2011, pp. 327–336.
- [15]. K. Lu, X. Zhou, X. Wang, W. Zhang, and G. Li, "RaceFree : An Efficient Multi-Threading Model for Determinism," in *In Proceedings of the 18th ACM SIGPLAN symposium on Principles and practice of parallel programming (PPoPP '13)*, 2013, pp. 297–298.
- [16]. R. Crawfis, "'Modern GPU Architecture', Lecture Notes for CSE 694G Game Design and Project." Ohio State University, Spring 2007.
- [17]. E. Lindholm, J. Nickolls, S. Oberman, and J. Montrym, "NVIDIA Tesla: A unified graphics and computing architecture.," *IEEE Microarchitecture*, vol. 28, no. 2, pp. 39–55, 2008.
- [18]. B. J. D. Owens, M. Houston, D. Luebke, S. Green, J. E. Stone, and J. C. Phillips, "GPU computing," *Proceedings of the IEEE* 96, no. 5, pp. 879–899, 2008.
- [19]. NVIDIA, "CUDA C Programming Guide," no. July. NVIDIA Corporation, 2013.
- [20]. Wikipedia, "General-purpose computing on graphics processing units," 2013. [Online]. Available: <http://tiny.cc/g4pw6w>. [Accessed: 21-Nov-2013].
- [21]. Wikipedia, "CUDA," *Wikipedia, The Free Encyclopedia*, 2013. [Online]. Available: <http://en.wikipedia.org/wiki/CUDA>. [Accessed: 22-Nov-2013].
- [22]. C. Woolley, "CUDA Overview Tutorial." Developer Technology Group, NVIDIA Corporation. Available Online: <http://www.cc.gatech.edu/~vetter/keeneland/tutorial-2011-04-14/02-cuda-overview.pdf>
- [23]. J. Nickolls, I. Buck, M. Garland, and K. Skadron, "Scalable Parallel with CUDA," *Queue*, vol. 6, no. 2, pp. 40–53, 2008.
- [24]. N. Gupta, "CUDA Programming, Complete Reference on CUDA," *CUDA programming Blog*, 2013. [Online]. Available: <http://tiny.cc/2niy6w>. [Accessed: 22-Nov-2013].
- [25]. OLCF, "Oak Ridge National Labs Accelerated Computing Guide: CUDA Vector Addition," *Accelerator Programming*, 2013. [Online]. Available: <https://www.olcf.ornl.gov/tutorials/cuda-vector-addition/>. [Accessed: 24-Nov-2013].
- [26]. GWU, "Supercomputers - HPCL – The George Washington University High Performance Computing Lab.," 2013. [Online]. Available: <http://hpcl2.hpcl.gwu.edu/index.php/supercomputers> . [Accessed: 25-Nov-2013].
- [27]. B. Barney, "POSIX threads programming." Lawrence Livermore National Laboratory, 2009. [Online]. Available: <https://computing.llnl.gov/tutorials/pthreads/> . [Accessed: 11-Dec-2013].
- [28]. Opengroup.org, "A Backgrounder on IEEE Std 1003.1, 2003 Edition," POSIX: A Backgrounder on IEEE Std 1003.1, 2003 Edition, 2003. [Online]. Available: <http://www.opengroup.org/austin/papers/backgrounder.html>. [Accessed: 11-Dec-2013].
- [29]. Wikipedia, "Pthreads," Pthreads, 2013. [Online]. Available: [http://en.wikipedia.org/wiki/POSIX\\_Threads](http://en.wikipedia.org/wiki/POSIX_Threads). [Accessed: 11-Dec-2013].
- [30]. A. Nin, J. Diaz, and C. Mun, "A Survey of Parallel Programming Models and Tools in the Multi and Many-Core Era," *IEEE Trans. Parallel Distrib. Syst.*, vol. 23, no. 8, pp. 1369–1386., 2012.
- [31]. J. Cook, "Pthreads, Tutorials and Tools," Tutorials and Tools, Computer Science Department, New Mexico University. [Online]. Available: [http://www.cs.nmsu.edu/~jcook/Tools/pthreads/pthread\\_con d.html](http://www.cs.nmsu.edu/~jcook/Tools/pthreads/pthread_con d.html). [Accessed: 11-Dec-2013].
- [32]. F. Mueller, "A Library Implementation of POSIX Threads under UNIX," in *In USENIX Winter*, 1993, pp. 29–42.
- [33]. J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, pp. 107-113, 2008.
- [34]. F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, *et al.*, "Bigtable: A distributed storage system for structured data," *ACM Transactions on Computer Systems (TOCS)*, vol. 26, p. 4, 2008.
- [35]. S. Ghemawat, H. Gobioff, and S.-T. Leung, "The Google file system," in *ACM SIGOPS Operating Systems Review*, 2003, pp. 29-43.