

Muhammad Abdelghaffar Awad

CONTACT INFORMATION	<i>E-mail:</i> mawad@ucdavis.edu <i>GitHub:</i> github.com/maawad	<i>Website:</i> maawad.github.io
EDUCATION	University of California, Davis <i>Electrical and Computer Engineering Department</i> Ph.D., Electrical and Computer Engineering, June 2022 M.S., Electrical and Computer Engineering, June 2022 Advisor: Professor John D. Owens Dissertation Topic: “Fully Concurrent GPU Data Structures” Alexandria University <i>Naval Architecture and Marine Engineering Department</i> B.S., with honours, Naval Architecture and Marine Engineering, June 2013	Davis, California 2016–2022 Alexandria, Egypt 2009–2013
PROFESSIONAL EXPERIENCE	AMD Research Member of Technical Staff University of California, Davis Electrical and Computer Engineering Department Graduate Student Researcher Lead research and open-source development of novel, state-of-the-art concurrent dynamic GPU data structures using CUDA and C++. With my collaborators, I published four first author publications where I was responsible for the paper writeup, managing and assigning tasks to coauthors, and designing and implementing the following data structures (available on GitHub): <ul style="list-style-type: none">– Dynamic GPU B-Tree– Multiversion GPU B-Tree– Dynamic hash-based GPU graph data structure (integrated into Gunrock)– GPU (on-device) epoch-based memory reclamation scheme to safely reclaim memory for concurrent GPU data structures– Static GPU hash tables that use probing schemes including cuckoo hashing, power-of-two choices, and iceberg hashing– Bucketed cuckoo hash set written in PTX and JIT-compiled. NVIDIA Research Intern Designed and implemented dynamic GPU string data structures to efficiently store and manipulate variable-sized strings. Activision Programming Intern Implemented and tested foliage rendering algorithms in HLSL shaders to produce high-quality images efficiently. Worked with team members to build a QT-based GUI tool to procedurally generate tree models.	Santa Clara, California August 2022– Davis, California September 2016–June 2022 Santa Clara, California June–September 2020 Redmond, Washington July–September 2017
TEACHING EXPERIENCE	University of California, Davis Electrical and Computer Engineering Department Teaching Assistant Course: Engineering Problem Solving (ENG 6)	Davis, California September–December 2017

Held office hours and managed lab sessions where I taught students how to solve problems and debug code more effectively. Created assignments to cover a wide range of topics and improve students' understanding. Graded assignments.

**Arab Academy for Science, Technology
and Maritime Transport**

Alexandria, Egypt

College of Maritime Transport and Technology

Teaching Assistant

July 2014–August 2016

Courses: Ship Design (MM543T) and Naval Architecture (MM241T)

Helped professors with lecture preparation, teaching, and general class management. Improved students' understanding through carefully constructed assignments and exams. Graded assignments and exams.

REFEREED
PUBLICATIONS

1. **M. A. Awad**, S. Ashkiani, S. D. Porumbescu, M. Farach-Colton and J. D. Owens. "Analyzing and Implementing GPU Hash Tables." *SIAM Symposium on Algorithmic Principles of Computer Systems, APOCS 2023*. January 2023.
2. **M. A. Awad**, S. D. Porumbescu and J. D. Owens. "A GPU Multiversion B-Tree." *Proceedings of the 31st International Conference on Parallel Architectures and Compilation Techniques, PACT 2022*. October 2022.
3. **M. A. Awad**, S. Ashkiani, S. D. Porumbescu and J. D. Owens. "Dynamic Graphs on the GPU." *Proceedings of the 34th IEEE International Parallel and Distributed Processing Symposium, IPDPS 2020*. May 2020.
4. **M. A. Awad**, S. Ashkiani, R. Johnson, M. Farach-Colton and J. D. Owens. "Engineering a High-Performance GPU B-Tree." *Proceedings of the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*. February 2019.
5. S. A. Mitchell, M. S. Ebeida, **M. A. Awad**, C. Park, A. Patney, and A. Rushdi. "Spoke-Darts for High-Dimensional Blue-Noise Sampling." *ACM Transactions on Graphics (TOG)*. July 2018.
6. **M. A. Awad**, A. Rushdi, M. A. Abbas, S. A. Mitchell, A. H. Mahmoud, C. L. Bajaj, M. S. Ebeida. "All-Hex Meshing of Multiple-Region Domains without Cleanup." *Proceedings 25th International Meshing Roundtable (IMR25)*. September 2016.
7. M. S. Ebeida, A. Rushdi, **M. A. Awad**, A. H. Mahmoud, D.-M. Yan, S. English, J. D. Owens, C. Bajaj, and S. A. Mitchell. "Disk Density Tuning of a Maximal Random Packing." *Proceedings of the Symposium on Geometry Processing (SGP 2016)*. June 2016.
8. M. S. Ebeida, S. A. Mitchell, A. Patney, A. A. Davidson, S. Tzeng, **M. A. Awad**, A. H. Mahmoud, and J. D. Owens. "Exercises in High-Dimensional Sampling: Maximal Poisson-disk Sampling and k-d Darts." In Janine Bennett, Fabien Vivodtzev, and Valerio Pascucci, editors, *Topological and Statistical Methods for Complex Data Tackling Large-Scale, High-Dimensional, and Multivariate Data Sets*, Springer. November 2014.
9. M. S. Ebeida, **M. A. Awad**, X. Ge, A. H. Mahmoud, S. A. Mitchell, P. M. Knupp, and L.-Y. Wei. "Improving Spatial Coverage while Preserving Blue Noise of Point Sets." *Computer Aided Design (SIAM GD/SPM 2013)*. November 2013.
10. M. S. Ebeida, A. H. Mahmoud, **M. A. Awad**, M. A. Mohammed, S. A. Mitchell, A. Rand, and J. D. Owens. "Sifted Disks." *Computer Graphics Forum (Eurographics 2013)*. May 2013.

OTHER PUBLICATIONS	<ol style="list-style-type: none"> 1. M. A. Awad, S. Ashkiani, S. D. Porumbescu, M. Farach-Colton and J. D. Owens. “Better GPU Hash Tables.” <i>CoRR</i>, abs/2108.07232(2108.07232). August 2021. 2. S. A. Mitchell, M. A. Awad, M. S. Ebeida and Laura P. Swiler. “Fast Approximate Union Volume in High Dimensions with Line Samples.” Technical Report, <i>Sandia National Laboratories</i>. August 2018.
TALKS	<p>A GPU Multiversion B-Tree.</p> <ul style="list-style-type: none"> – 31st International Conference on Parallel Architectures and Compilation Techniques (PACT), Chicago, October 2022. <p>Better GPU Hash Tables.</p> <ul style="list-style-type: none"> – CUDA Community Meetup Group, Virtual. October 2021. <p>Dynamic Graphs on the GPU.</p> <ul style="list-style-type: none"> – 34th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2020), Virtual. May 2020. <p>Engineering a High-Performance GPU B-Tree.</p> <ul style="list-style-type: none"> – Second Hawaii Workshop on Parallel Algorithms and Data Structures, University of Hawaii at Manoa, Hawaii. December 2019. – NVIDIA, Santa Clara, California. April 2019. – 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2019), Washington, DC. February 2019. <p>All-Hex Meshing of Multiple-Region Domains without Cleanup.</p> <ul style="list-style-type: none"> – 25th International Meshing Roundtable (IMR25), Washington, DC. September 2016.
SERVICE	Reviewer for IEEE Transactions on Parallel and Distributed Systems (TPDS), 2019
AWARDS	Dissertation Fellowship, University of California, Davis, Department of Electrical and Computer Engineering, Spring Quarter 2022
TECHNICAL SKILLS	<p>Programming: C++, HIP, CUDA C/C++, NVIDIA’s PTX, QT, OpenGL.</p> <p>Applications: NVIDIA Nsight Compute and Systems, Microsoft Visual Studio, AutoCAD, ParaView, L^AT_EX.</p> <p>Operating Systems: Microsoft Windows, Linux.</p> <p>Other Tools: Git, CMake, Visual Studio Code.</p>