

John Douglas Owens

Electrical and Computer Engineering, University of California
One Shields Avenue, Davis, CA 95616-5294 USA
+1-530-754-4289 / jowens@ece.ucdavis.edu
<http://www.ece.ucdavis.edu/~jowens/>

RESEARCH INTERESTS	Computer systems; parallel computing, general-purpose computation on graphics hardware / GPU computing, parallel algorithms and data structures, graphics architectures, parallel architectures and programming models.
EDUCATION	<p>Stanford University Stanford, California <i>Department of Electrical Engineering</i> 1995–2003 Ph.D., Electrical Engineering, January 2003 M.S., Electrical Engineering, March 1997 Advisors: Professors William J. Dally and Pat Hanrahan Dissertation Topic: “Computer Graphics on a Stream Architecture”</p> <p>University of California, Berkeley Berkeley, California <i>Department of Electrical Engineering and Computer Sciences</i> 1991–1995 B.S., Highest Honors, Electrical Engineering and Computer Sciences, June 1995</p>
PROFESSIONAL APPOINTMENTS	<p>University of California, Davis Davis, California Child Family Professor of Engineering and Entrepreneurship 2014– Associate Professor 2008–2014 Assistant Professor 2003–2008 Professor in the Department of Electrical and Computer Engineering; member of Graduate Groups of Electrical and Computer Engineering and Computer Science. Visiting Scientist, Lawrence Berkeley National Laboratory.</p> <p>Google San Francisco, California Visiting Researcher (Sabbatical) August 2024–August 2025 Incubating portable compute shader algorithms for the WebGPU API.</p> <p>Twitter San Francisco, California Software Engineer (Sabbatical) July–December 2012 Engineer in Core System Libraries group in Twitter’s Runtime Systems division.</p> <p>Stanford University Stanford, California Research Assistant 1997–2002 An architect of the Imagine Stream Processor under the direction of Professor William J. Dally. Responsible for major portions of hardware and software design for Imagine and its tools and applications.</p> <p>Stanford University Stanford, California Teaching Assistant Fall 2001 Teaching assistant for Computer Science 99s, “The Coming Revolution in Computer Architecture”, under Professor William J. Dally. Designed course with Professor Dally, including lecture topics, readings, and laboratories.</p> <p>Interval Research Corporation Palo Alto, California Consultant 1997–2000 Investigated new graphics architectures under the direction of Dr. Matt Regan.</p> <p>Stanford University Stanford, California Research Assistant 1995–1997 Under the direction of Professor Pat Hanrahan, designed and built the Lightning distributed framebuffer.</p> <p>Silicon Studio, Silicon Graphics Inc. Mountain View, California</p>

Software Engineer Summer 1995
 Performed development work on SGI's Firewalker video-game authoring system, mastering game titles to game platforms.

University of California, Berkeley Berkeley, California
 Teaching Assistant Spring 1995
 Teaching assistant for Computer Science 150, "Digital Design", under Professor Richard Newton. Responsible for laboratory section, office hours, grading, and midterm review.

Intel Corporation, P7 Architecture Group Santa Clara, California
 Design Engineer Summer 1994
 Designed and implemented graphical user interface to Intel's Sphinx microarchitecture simulator.

First Person Inc. (subsidiary of Sun Microsystems Inc.) Palo Alto, California
 Hardware Designer Summer 1993
 Assisted in design of NTSC SBus-compatible framebuffer.

ViewStar Corporation Emeryville, California
 Quality Assurance Engineer Summer 1992
 Tested and debugged ViewStar's Document Management System.

Oracle Corporation Redwood Shores, California
 Technical Staff Summer 1991
 Tested and debugged Oracle for Macintosh.

AWARDS AND HONORS

Editors' Pick for Notable Papers, ACM Transactions on Mathematical Software (C. Yang, A. Buluç, J. D. Owens, *GraphBLAST: A High-Performance Linear Algebra-based Graph Framework on the GPU*). Awarded December 2024

UC Davis College of Engineering Outstanding Mid-Career Faculty for Excellence in Research Award, 2022

IEEE Fellow, 2021

American Association for the Advancement of Science (AAAS) Fellow, 2021

Department of Electrical and Computer Engineering Distinguished Research Award, 2019 (first recipient)

Test of Time Award, 2019 High Performance Graphics, for the most influential paper from the 2010 predecessor conferences to HPG (S. Tzeng, A. Patney, and J. D. Owens, *Task Management for Irregular-Parallel Workloads on the GPU*)

Principal Investigator, NVIDIA GPU AI Research Center at UC Davis, 2019–present

MIT/IEEE/Amazon GraphChallenge:

- 2019 Student Innovation Award (X. Wang, Z. Lin, C. Yang, and J. D. Owens, *Accelerating DNN Inference with GraphBLAS and the GPU*)
- 2019 Finalist (L. Wang and J. D. Owens, *Fast BFS-Based Triangle Counting on GPUs*)

Distinguished Paper, European Conference on Parallel and Distributed Computing (Euro-Par) 2018 (C. Yang, A. Buluç, and J. D. Owens, *Design Principles for Sparse Matrix Multiplication on the GPU*)

Best Artifact Award, European Conference on Parallel and Distributed Computing (Euro-Par) 2018 (C. Yang, A. Buluç, and J. D. Owens, *Design Principles for Sparse Matrix Multiplication on the GPU*)

"Advisor of Excellence", Theta Tau Professional Engineering Fraternity, 2018, Omicron Gamma chapter (national award, selected from advisors of 83 chapters)

ACM Distinguished Member, 2017

Test of Time Award, 2017 High Performance Graphics, for the most influential paper from the 2007 and 2008 predecessor conferences to HPG (S. Sengupta, M. Harris, Y. Zhang, and J. D. Owens, *Scan Primitives for GPU Computing*)

Distinguished Paper, ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2016 (Y. Wang, A. Davidson, Y. Pan, Y. Wu, A. Riffel, and J. D. Owens, *Gunrock: A High-Performance Graph Processing Library on the GPU*)

Distinguished Paper, European Conference on Parallel and Distributed Computing (Euro-Par) 2015 (L. Wang, S. Baxter, and J. D. Owens, *Fast Suffix Array on the GPU*)

Best Paper Finalist, 2015 IEEE International Symposium on Workload Characterization (Y. Wu, Y. Wang, Y. Pan, C. Yang, and J. D. Owens, *Performance Characterization of High-Level Programming Models for GPU Graph Analytics*)

US Frontiers of Engineering Symposium (National Academy of Engineering), 2016, 2013

NVIDIA CUDA Fellow, 2012

“Advisor of Excellence”, Theta Tau Professional Engineering Fraternity, 2012, Omicron Gamma chapter (national award, selected from advisors of 55 chapters)

Finalist, ASUCD Excellence in Education Award (College of Engineering), 2009

ECE Graduate Student Association Award for Graduate Teaching and Mentorship, 2009 (first recipient)

Best Paper Award, Graphics Hardware 2007 (S. Sengupta, M. Harris, Y. Zhang, and J. D. Owens, *Scan Primitives for GPU Computing*)

NVIDIA Faculty Teaching Fellowship, 2006

Department of Energy Early Career Principal Investigator Award, 2004

Stanford Program for Academic Excellence Mentor, 1997–2000

Stanford College of Engineering Lawrence R. Thielen Memorial Fellowship

Charles Mills Gayley Fellowship for Graduate Study

Eta Kappa Nu (Mu Chapter)

Tau Beta Pi (California Alpha Chapter)

**PROFESSIONAL
SERVICE**

Leadership Service

Chair, OpenCilk Academic Board (2022–present)

Emerging Models of Colossal Computation ($E = mc^2$)

- Organizing Committee (2023)

US Frontiers of Engineering Symposium

- Organizing Committee (2016)

PeerJ Computer Science

- Editorial board (2015–present)

Innovative Parallel Computing (InPar)

- Paper Chair (2012)

Siggraph/Eurographics High Performance Graphics

- Steering Committee (2011–2024)
- General Chair (2011)
- Program Chair (2009)

Graphics Hardware

- Papers Chair (2008)
- Publicity Chair (2005, 2006, 2007)

IS&T/SPIE Electronic Imaging: Parallel Processing for Imaging Applications

- Conference Chair (2011)

2006 Workshop on On- and Off-Chip Interconnection Networks for Multicore Systems

- Chair

Program/Paper Committee Service

IEEE International Parallel & Distributed Processing Symposium (IPDPS)
— Program Committee (2024, 2023, 2016)

IEEE International Conference on High Performance Computing (HiPC)
— Program Committee (2023, 2022, 2021, 2018, 2015)

Siggraph/Eurographics High Performance Graphics
— Paper Committee (2024, 2023, 2021, 2020, 2019, 2018, 2017, 2016, 2015, 2014, 2013, 2012, 2010)

IEEE International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoc) / Special Session on Auto-Tuning for Multicore and GPU (ATMG)
— Program Committee (2024, 2023, 2022, 2021, 2020, 2019)

IEEE International Conference on Cluster Computing
— Program Committee (2022, 2017)

Workshop on Graphs, Architectures, Programming, and Learning (GrAPL)
— Program Committee (2019)

Tenth International Workshop on Accelerating Analytics and Data Management Systems Using Modern Processor and Storage Architectures (ADMS 2019)
— Program Committee (2019)

European Symposium on Algorithms
— Program Committee (Track B, Engineering and Applications Track, 2019)

Hawaiian Workshop on Parallel Algorithms and Data Structures
— Organizing Committee (2019, 2017)

Graph Algorithms Building Blocks
— Program Committee (2018, 2017, 2016)

Supercomputing Asia
— Technical Program Committee (2018)

Workshop on the Intersection of Graph Algorithms and Machine Learning (GraML)
— Program Committee (2018, 2017)

ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games
— Program Committee (2017, 2009, 2008)

High Performance Graph Data Management and Processing Workshop (HPGDMP)
— Program Committee (2016)

ACM/SIGARCH International Conference on Supercomputing (ICS)
— Program Committee (2015)

Eurographics Symposium on Rendering
— International Program Committee (2015)

International Conference on Multicore Software Engineering, Performance, and Tools (MUSEPAT)
— Program Committee (2015, 2013)

ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)
— External Review Committee (ERC) (2014)

1st International Workshop on OpenCL
— Program Committee (2013)

Heterogeneous and Unconventional Cluster Architectures and Applications Workshop
— Program Committee (2013)

ADBIS Workshop on GPUs in Databases (GID)
— Program Committee (2013, 2012, 2011)

IEEE International Conference on High Performance Computing (HiPC) Student Research Symposium
— Program Committee (2013)

ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming (PPoPP)
 — Program Committee (2012)

International Conference on Parallel Processing (ICPP)
 — Program Committee (2012)

20th Euromicro International Conference on Parallel, Distributed and Network-based Processing,
 Special Session on GPU Computing and Hybrid Computing
 — Program Committee (2012)

Symposium on Application Accelerators in High Performance Computing
 — Program Committee (2012, 2011, 2009)

Second Workshop on Hybrid Multi-core Computing (WHMC 2011)
 — Program Committee (2011)

First International Workshop on Characterizing Applications for Heterogeneous Exascale Sys-
 tems (CACHES)
 — Program Committee (2011)

IEEE/ACM Supercomputing
 — Posters Committee (2011)
 — Program Committee (Applications) (2009, 2008)

Eurographics
 — Program Committee, “Computational Graphics” (2009)
 — International Program Committee, Short Papers (2007)

Workshop on General Purpose Processing on Graphics Processing Units
 — Program Committee (Second Workshop, 2009)
 — Program Committee (First Workshop, 2007)

2nd Workshop on Programming Models for Ubiquitous Parallelism (PMUP 2007)
 — Program Committee

External Participant (report contributor), Recommendations from the DOE/ASCR Workshop on
 Visual Analysis and Data Exploration for the Exascale Era (2007)

GPGPU Workshop at the 6th International Conference on Computational Science (ICCS 2006)
 — Program Committee

Graphics Hardware
 — Program Committee (2004)

International Workshop on Networked Sensing Systems
 — Program Committee (2005, 2004)

Government Review Service

Department of Energy Exascale Computing Project, Co-Design Centers for Exascale Applications,
 reviewer

Department of Energy Early Career Research Program, reviewer

Los Alamos National Laboratory Information & Knowledge Sciences (IKS) Capability Review
 Committee (2011)

Department of Energy’s Small Business Innovative Research Program, panelist

National Science Foundation proposal panelist and reviewer

**UNIVERSITY
 SERVICE**

Graduate Group of Computer Science, Graduate Faculty Advisor, 2023–24, 2025–26

Working Group on Ensuring Freedom of Expression at UC Davis, 2017

College of Engineering Strategic Planning Committee, UC Davis, 2017

UC Davis Academic Senate Representative Assembly
 — Department representative, 2013–15, 2015–17
 — Department alternate, 2011–13

Dean of Engineering Search Advisory Committee, UC Davis, 2014–15
Chair, Department of Electrical and Computer Engineering Industrial Affiliates Committee, 2011–12, 2012–13
Athletics Administrative Advisory Committee (as representative of the Academic Senate), UC Davis, 2005–06, 2006–07, 2007–08, 2008–09, 2009–10, 2010–11
— Financial Aid Task Force (subcommittee of AAAC), 2006–2007
— Street Drug Task Force (subcommittee of AAAC), 2007–2008
— Budget Committee (subcommittee of AAAC), 2010–2011
Committee on Academic Support in Athletics, UC Davis, 2006–2007
Faculty advisor, Theta Tau Engineering Fraternity, UC Davis, 2005–

PATENTS

US 7589719. “Fast Multi-pass Partitioning via Priority Based Scheduling”. John Douglas Owens, Andy Riffel, Aaron Lefohn, Mark Leone, and Kiril Vidimce. Issued 15 September 2009.
US 7100026: “System and Method for Performing Efficient Conditional Vector Operations for Data Parallel Architectures.” William J. Dally, Scott Rixner, John D. Owens, and Ujval J. Kapasi. Issued 29 August 2006.
US 6269435: “System and method for implementing conditional vector operations in which an input vector containing multiple operands to be used in conditional operations is divided into two or more output vectors based on a condition vector”. William J. Dally, Scott Whitney Rixner, John Owens, and Ujval J. Kapasi. Issued 31 July 2001.

GRANT/GIFT FUNDING

Advanced Micro Devices (AMD). *A Library for Road Network Optimization on AMD APUs*. Unrestricted gift. Jan. 2025–Jun. 2025 (\$50,000).
National Science Foundation (Award # CCF-2403389). *SHF: Medium: A Tensor Formulation for Graph Analytics*. PI J. Owens. 1 Sept. 2024–31 Aug. 2028. \$1,091,988.
National Science Foundation (Award # TI-2346223). *NSF POSE: Phase II: Open-Source Ecosystem for OpenCilk*. PI C. E. Leiserson (MIT), co-PIs J. Owens, Y. Sun (UCR). 1 Aug. 2024–31 Jul. 2026. \$1,500,000.
Advanced Micro Devices (AMD). *GPU Kernels for Block-Sparse Multi-Head Attention*. Unrestricted gift. Jan. 2024–Dec. 2024 (\$75,000), Jan. 2025–December 2025 (\$100,000),
MLCommons Association. *Automotive Computing Research Fund*. Unrestricted gift. January 2023 (\$60,000), renewed March 2024 (\$90,000), June 2025 (\$45,000), December 2025 (\$45,000).
National Science Foundation (Award # TI-2229704). *POSE Phase I: Open Source Ecosystem for OpenCilk*. PI C. E. Leiserson (MIT). 15 Sept. 2022–31 Aug. 2023. \$300,000 (total), \$20,000 (UCD).
Advanced Micro Devices (AMD). *A Unified Framework for Trace Collection, Analysis, and Translation of GPU Workloads*. Unrestricted gift. Jan. 2022–Dec. 2022 (\$100,000), renewed Jan. 2023–Dec. 2023 (\$100,000).
US Geological Survey (Grant # G21AP10845-00). *HEC-RAS 2D Graphics Processor Unit (GPU) Solver*. PI J. Owens. 1 October 2021–30 September 2023 (\$329,611), extended to 30 September 2025 (\$213,400).
Computing Research Association. *Computing Innovation Fellows 2021 Project*. PI J. Owens. 1 Jan. 2022–31 May 2024. \$301,508.
Facebook. Unrestricted gift. Jan. 2021–Dec. 2022. \$240,000.
Semiconductor Research Corporation (Contract # 2019-AU-2927). *Benchmarking Automotive Workloads for Autonomous Driving*. PI J. Owens. 1 Jul. 2020–31 Jun. 2023. \$394,767.
NVIDIA AI Laboratory. *UC Davis Center for GPU Graph Analytics*. Unrestricted gift + NVIDIA DGX Station. Mar. 2019. \$149,000. Renewed Mar. 2020 (\$100,000), May 2022 ((\$100,000).
National Science Foundation (Award # CCF-1823037). *SPX: Collaborative Research: Global Address Programming with Accelerators*. PI J. Owens, co-PIs K. Yelick, A. Buluç (UC Berkeley).

1 Oct. 2018–30 Sept. 2021 (no-cost extension to 30 Sept. 2022). \$386,025 (UC Davis).
Department of Defense Advanced Research Projects Agency (DARPA). *SYMPHONY: Orchestrating Sparse and Dense Data for Efficient Computation*. Award # HR0011-18-3-0007. 1 June 2018–30 Oct 2022. \$827,790 (UC Davis).

Department of Defense Advanced Research Projects Agency (DARPA). *A Commodity Performance Baseline for HIVE Graph Applications*. Award # FA8650-18-2-7835. 8 June 2018–7 Sept. 2019, \$500,000; 30 Nov. 2019–31 May 2021, \$750,000.

Department of Defense Advanced Research Projects Agency (DARPA). *A Commodity Performance Baseline for HIVE Graph Applications*. Northrup Grumman, prime contractor. Nov. 2017–Apr. 2018. \$87,500. 28 Jan.–26 Feb. 2019. \$107,870.

Intel Corporation. *A Heterogeneous Programming Model for Graphics* (Award # A18-0805). 1 Oct. 2017–31 Dec. 2019. \$79,532.

Activision. Unrestricted gift. 30 Oct. 2017. \$15,000.

Adobe Data Science Research Award. *Scalability and Mutability for Large Streaming Graph Problems on the GPU*. PI J. Owens. Aug. 2017. \$50,000.

National Science Foundation (Award # OAC-1740333) *SI2-SSE: Gunrock: High-Performance GPU Graph Analytics*. PI J. Owens. 1 Oct. 2017–30 Sept. 2020. \$400,000.

UC Davis New Research Initiatives and Collaborative Interdisciplinary Research Grant. *Implementing Sparse Tensors on GPUs*. 27 July 2017–30 Sept. 2018. \$14,491.

Adobe Systems Inc. Unrestricted gift. Jan. 2017. \$7,500.

Intel Corporation. *Platform Assessment for Autonomous Vehicles*
— Award # 201604270: 1 Oct. 2016–31 Mar. 2021, \$150,000.
— Award # A18-0530: 1 Aug. 2017–30 June 2021. \$150,000.
— Award # A19-1613: 1 July 2018–31 Dec. 2021. \$100,000.

Sandia National Laboratories. *Mesh Improvement and Generation Algorithms*. 21 Sept. 2016–30 Sept. 2017. \$43,750.

National Science Foundation (Award # DMS-1622501) *High-Performance, High-Level Tools for Statistical Inference and Unsupervised Learning*. Program: Computational and Data-Enabled Science and Engineering in Mathematical and Statistical Sciences. PI J. Owens; co-PIs J. Bezanson, J. Fisher, A. Edelman (MIT). 15 Sept. 2016–31 Dec. 2019 (NCE until Dec. 2020). \$480,000 (total), \$164,612 (UCD).

National Science Foundation (Award # CCF-1637442) *Theory and implementation of dynamic data structures for the GPU*. Program: AitF—Algorithms in the Field. PI J. Owens; co-PI M. Farach-Colton (Rutgers). Sept. 2016–Aug. 2021. \$438,876 (UCD).

National Science Foundation (Award # CCF-1629657) *PARAGRAPH: Parallel, Scalable Graph Analytics*. XPS—Exploiting Parallelism & Scalability. PI P. Sadayappan (Ohio State); co-PIs J. Owens, S. Parthasarathy (Ohio State), L.-N. Pouchet (Ohio State). Sept. 2016–Aug. 2019 (NCE until Aug. 2021). Proposed \$375,000 (UCD); funded at \$328,123.

Sandia National Laboratories. *Scalable Methods for High-Dimensional Numerical Integration and Uncertainty Quantification Problems*. 1 Oct. 2015–30 Sept. 2016. \$75,000.

Department of Defense Advanced Research Projects Agency (DARPA) SBIR SB152-004. *Many-Core Acceleration of Common Graph Programming Frameworks*. PI G. Jayachandran (Onu Technology, Inc.), co-PI J. Owens. Funded directly as Phase II: award W911NF-16-C-0020. Jan. 2016–Dec. 2017. Total funding: \$1,399,583.

Stratovan Corporation. Unrestricted gift. Sept. 2015–Jun. 2022. \$283,540.58.

Adobe Digital Marketing Research Award. *Scaling Convex Optimization with GPUs*. PI J. Owens, co-PI S. Boyd (Stanford University). Oct. 2014. \$50,000.

National Academy of Engineering Frontiers of Engineering Program, based on funding from The Grainger Foundation. *Parallel Matrix Factorization: Towards GPUs in the Data Center*. PI

J. Owens, co-PI T. Sahai (United Technologies Research Center). 1 April 2014–31 May 2016. \$30,000.

Department of Defense Advanced Research Projects Agency (DARPA) STTR ST13B-004 (“Data-Parallel Analytics on Graphics Processing Units (GPUs)”). *A High-Level Operator Abstraction for GPU Graph Analytics*. PI E. Elsen (Royal Caliber LLC), co-PI J. Owens. Phase I: Jan.-June 2014, award D14PC00023. \$70,000 (UCD: \$37,500). Phase II: 15 Jan. 2015–31 March 2016, award D15PC00010. (UCD: \$370,000).

Department of Defense (XDATA program). *An XDATA Architecture for Federated Graph Models and Multi-Tier Asymmetric Computing*. Oct. 2012–Sept. 2017. Prime contractor: Sotera Defense Solutions, Inc., US Army award W911QX-12-C-0059 (prime contractor Oct. 2012–July 2013: SYSTAP LLC; prime contractor Aug. 2013–July 2016: L-3 Data Tactics). UCD funding: Oct. 2012–July 2013: \$74,230.37; Aug. 2013–Dec. 2013: \$119,443; Jan. 2014–Jan. 2016: \$629,203; Aug. 2016–Sept. 2017: \$400,000.

Department of Energy, Sandia National Laboratories, Laboratory Directed Research and Development (LDRD) program. *Efficient Probability of Failure Calculations for QMU using Computational Geometry*. (Proposal No. 13-0144, award #1292969). PI S. Mitchell (Sandia National Laboratories), co-PIs M. Ebeida, B. Adams, V. Romero, J. Owens. FY13–15 (22 Oct. 2012–31 Oct. 2014). \$1.6M (UCD portion \$120,000).

UC Lab Fees Research Program Award 12-LR-238449. *Probabilistic Algorithms for New Computer Architectures*. PI L. Monroe (Los Alamos National Laboratory), co-PIs J. Wendelberger (LANL), S. Michalak (LANL), J. Owens. July 2012–Sept. 2015 (no-cost extension to Sept. 2016). \$764,396.

Advanced Micro Devices (AMD). Unrestricted gift. June 2012. \$75,000.

National Science Foundation Supplemental Funding Request (high-school-student research supplement), *SHF:Small:Software Fundamentals for Manycore Systems* (award # 1017399, existing award # 1017399). PI J. Owens. Awarded June 2012. \$4800.

Intel Science and Technology Center for Visual Computing. PI P. Hanrahan (Stanford), theme leads M. Agrawala (Berkeley), D. James (Cornell), J. Owens (UC Davis), S. Seitz (Washington). 8 universities participating, ~30 PIs. Jan. 2011–Dec. 2015. \$15M over 5 years.

National Science Foundation (Award # OCI-1032859) *SDCI: HPC: Improvement: Infrastructure for Multi-Node Manycore Computing*. Office of Cyberinfrastructure. PI J. Owens. Sept. 2010–Aug. 2015. \$391,859.

National Science Foundation (Award # CCF-1017399) *SHF:Small:Software Fundamentals for Manycore Systems*. Division of Computer and Communication Foundations—Software and Hardware Foundations. PI J. Owens. Aug. 2010–July 2015. \$499,825.

Center for Information Technology Research in the Interest of Society (CITRIS) seed funding grant *Computational Tools for River and Estuary Flow Prediction*. PI J. Owens, co-PI B. Younis. Fall 2010–Summer 2011. \$70,810.

National Science Foundation (Award # IIS-0964357) *HCC: Medium: Collaborative Research: Data-Parallel Hash Tables: Theory, Practice and Applications*. IIS—Human-centered Computing program. PI A. Amenta; co-PIs J. Owens, M. Mitzenmacher (Harvard). Aug. 2010–July 2013. Proposed \$677,994 (UCD); funded at \$532,084.

Department of Defense / Air Force Office of Scientific Research Small Business Technology Transfer (STTR) Program. *Innovative CFD Algorithms, Libraries and Python Frameworks for Hybrid CPU-GPU Compute Architectures* (STTR Topic Number F09B-T18). PI: Dr. E.P.N. Duque (JMSI, Inc. / Intelligent Light). University PIs: J. Owens, R. Davis. \$30,000 (UCD). 1 April–20 Dec. 2010.

Microsoft. Unrestricted gift. April 2010. \$30,000.

National Nuclear Security Administration (NNSA). *Real-Time Three-Dimensional Exploration of Wide-Area High-Resolution Aerial Video* (Award # SC09-PDP08-23). PI K. Joy, co-PI J. Owens. \$800,000.

HP Labs Innovation Research Program. *GPU Completeness: Computational Patterns and Analysis for GPU Computing*. \$70,457 in first year, \$70,000 for second year, \$18,800 for third year. Awarded in June 2009, renewed in May 2010 and May 2011.

UCD 2009–2010 Committee on Research Small Grant in Aid Award. May 2009. \$2,000.

Intel Microprocessor Technology Labs, *Speech Recognition on Commodity Processors: GPGPU vs. Larrabee*. Unrestricted gift. \$50,000. Aug. 2008–July 2009; renewed as grant (\$50,000), Sept. 2009–Sept. 2010.

UC MICRO proposal *Automotive Computing using the GPU* (award #08-57, tied to BMW, NVIDIA, and Rambus gifts dated March 2008). PI J. Owens. Sept. 2008–Dec. 2009. \$28,208.

BMW. Unrestricted gift. March 2008. \$25,000.

NVIDIA. Unrestricted gift. March 2008. \$25,000.

Rambus. Unrestricted gift. March 2008. \$15,000.

National Science Foundation Supplemental Funding Request (undergraduate research supplement), *Data Structures for Data-Parallel Architectures* (award # 0734164, existing award # 0541448). PI J. Owens. Awarded May 2007. \$6000.

UC Discovery Grant, Industry-University Cooperative Research Program, *Workshop on On- and Off-Chip Interconnection Networks for Multicore Systems* (award # COM06-324), PI J. Owens. \$15,000.

National Science Foundation (CCF), *Workshop on On-And Off-Chip Interconnection Networks for Multicore Systems* (award # 0644602), PI W. Dally, co-PI J. Owens, awarded Aug. 2006. \$29,638.

Undergraduate Education in Parallel Architecture, NVIDIA Teaching Fellowship, Fall 2006. \$25,000.

A Common Software Interface for GPU Data Structures, Los Alamos National Laboratory. PI J. Owens. 15 July 2006–15 Sept. 2006. \$32,663.

Department of Energy “Scientific Discovery through Advanced Computing” program. *SciDAC Institute for Ultra-Scale Visualization* (contract DE-FC02-06-ER25777). PI K.-L. Ma, co-PIs R. Ross, N. Max, H.-W. Shen, K. Moreland, J. Owens, J. Huang, G. Humphreys. 2006–2011. \$8,228,429.

UC MICRO proposal *Scientific Computing for Energy Technology Applications Using Graphics Hardware* (tied to Chevron gift dated Oct. 2005). PI J. Owens. Aug. 2006–Dec. 2007. \$16,625.

Intel Corporation Grant: *CEEL: Computer Engineering Education Laboratories with Wireless Networking Extension at UC Davis*. 2006. PI C.-N. Chuah; co-PIs S. Ghiasi, J. Owens, P. Mohapatra, K. Wilken. \$95,254 (\$50,000 from Intel).

Chevron. Unrestricted gift. Oct. 2005. \$20,000.

National Science Foundation (Award # CCF-0541448) *Data Structures for Data-Parallel Architectures*. Foundations of Computing Processes and Artifacts program. PI J. Owens. July 2006–June 2009. Proposed \$419,436; funded at \$200,000.

Lockheed-Martin (via DOD) *Assessment of GPUs for DoD DSP Applications*. PI J. Owens. June 2005–June 2006. \$50,000.

UC Davis–Los Alamos National Laboratory “Cooperative Agreement for Research and Education” (“CARE”) proposal *Solving the Multipass Partitioning Problem for Graphics Hardware*. PI J. Owens; co-PI P. McCormick, LANL. 1 Oct. 2004–31 Jan. 2005. \$54,708.

Chevron-Texaco. Unrestricted gift. Sept. 2004. \$20,000.

UC MICRO proposal *Scientific Computing for Energy Technology Applications Using Graphics Hardware* (tied to ChevronTexaco gift dated Dec. 2003). PI J. Owens. Aug. 2004–Dec. 2005. \$14,588.

Department of Energy Early Career Principal Investigator Award. *A Programming Framework for Scientific Applications on CPU-GPU Systems*. PI J. Owens. 3 years, 15 Aug. 2004–15 Aug. 2007, granted two one-year no-cost extensions. \$300,000.

Los Alamos National Laboratory. *GPU/MPP Investigations*. PI J. Owens, summer 2004. \$18,000.

Chevron-Texaco. Unrestricted gift. Dec. 2003. \$20,000.

UCD Faculty Research Grant Program. *The UC Davis Metanet: A Scalable, Ubiquitous, Extensible Framework for Sensor Networks*. PI J. Owens. Aug. 2003 (for 2003–2004 academic year). \$14,000.

UCD New Faculty Research Grant. *Client-Side Computation using Graphics Hardware for Video-conferencing and Immersive Environments*. PI J. Owens. Jan. 2003. \$4000.

EXTERNAL TALKS

- 2020 Apple Computer, NVIDIA GPU Technology Conference, UC Davis ECEXpo (Santa Clara), UC Berkeley (guest lecture, CS 267)
- 2019 Dagstuhl Seminar 19051, UC Berkeley (guest lecture, CS 267), Semiconductor Research Corporation/Semiconductor Industry Association/Department of Energy Decadal Plan Workshop on New Compute Trajectories for Energy-Efficient Computing
- 2018 Adobe, DARPA HIVE PI meeting, IEEE IPDPS, Symphony kickoff (MIT, Cambridge MA), DARPA SDH Symphony webinar, UC Davis EEC 298, DARPA HIVE project review (Arlington VA), NSF COGNAC project meeting, UC Davis EEC 1
- 2017 UC Berkeley (guest lecture, CS 267), NSF AiTF PI meeting, NVIDIA CRADA review (Graph and Data Analytics), NVIDIA GPU Graph Roundtable, AMD, Amazon
- 2016 Dagstuhl Seminar 16101, NVIDIA online webinar, National Academy of Engineering Frontiers of Engineering, UC Davis EEC 1 (Introduction to Electrical and Computer Engineering) guest lecture
- 2015 DARPA XDATA PI meeting, NVIDIA GPU Technology Conference Data Analytics Roundtable, Adobe Data Science Symposium, Indiana University CREST Reconfigurable Execution Framework Testbed Workshop, Lawrence Livermore National Laboratory, Netflix
- 2014 Dagstuhl Seminar 14091, Adobe Data Science Symposium, Twitter, United Technologies Research Center, Dagstuhl Seminar 14462
- 2013 Yale University, Intel Science and Technology Center for Visual Computing annual retreat, UC Davis Joint Senate/Administration Retreat on Online Education, UC Davis Department of Statistics seminar, UC Davis Summer Institute for Teaching with Technology, UC Davis EEC 1 (Introduction to Electrical and Computer Engineering) guest lecture
- 2012 Lawrence Livermore National Laboratory, Synchronization-reducing and Communication-reducing Algorithms and Programming Models for Large-scale Simulations Workshop (plenary talk), Intel Science and Technology Center for Visual Computing yearly retreat, Intel Science and Technology Center for Visual Computing review, Lawrence Berkeley National Laboratory, UC Davis Graduate Group of Applied Mathematics 7th Annual Mini-Conference, Advanced Micro Devices, Qualcomm Research Silicon Valley, Second Workshop on Architectures and Systems for Big Data (panel)
- 2011 Twitter, Northrup Grumman, Pacific Northwest National Laboratory, Harvard (guest lecture, CS 264), Intel Science and Technology Center for Visual Computing webinar, Intel Science and Technology Center for Visual Computing yearly retreat, Harvard (Computer Science Colloquium), MIT, Oxford University, Workshop on Hybrid Multi-core Computing, Infosys
- 2010 Dagstuhl Seminar 10091 (Data Structures), Stanford (CS 448s guest lecture), KLA Tenor, Lawrence Berkeley National Laboratory, Programming Environments for Emerging Parallel Systems (PEEPS), 6th International Workshop on Parallel Matrix Algorithms and Applications (PMAA '10), Oak Ridge National Laboratory, SciDAC 2010, Center for Scalable Application Development Software (CScADS) Workshop on Scientific Data and Analytics for Extreme Scale Computing, ARM External Research Speaker Conference
- 2009 Nara Institute of Science and Technology, UC Davis Faculty Mentoring Faculty Program, Intel Visual Computing Group Tech Summit (Folsom) (keynote), Heterogeneity in Computing Workshop (keynote), GPU Technology Conference (Santa Clara), Supercomputing 2009 (tutorial)
- 2008 Architectural Support for Programming Languages and Operating Systems (course), International Ph.D. School in Algorithms for Advanced Processor Architectures (Copenhagen), Lund University (Sweden), Workshop on Programming Massively Parallel Processors

- (Urbana-Champaign), ACM SIGGRAPH (course), HP Labs, IBM Watson, Cal Poly San Luis Obispo, UCSB
- 2007** Intel Microprocessor Research Lab (Santa Clara), International Workshop on Logic Synthesis, IEEE/NATEA Annual Conference (2007 New Frontiers in Computing Technology), Rambus Tech Forum (Santa Clara), ACM SIGGRAPH 2007 (course) (San Diego), Microsoft Research Asia, Beijing Capital Normal University, Baidu, Tsinghua University, Apple Computer, Supercomputing 2007 (Austin) (tutorial)
- 2006** Shell Research (Houston), UC Berkeley, Supercomputing 2006 (tutorial), Intel (Folsom)
- 2005** Microsoft Research, Intel (Santa Clara), National Center for Atmospheric Research, University of Colorado, Graphics Hardware (panel talk), Eurographics, IEEE Visualization (tutorial), Lawrence Livermore National Laboratory, Lockheed-Martin
- 2004** Microsoft Research, Intel Research (Berkeley), UCSF, High-Performance Embedded Computing Conference, IEEE Visualization (tutorial)
- 2003** Pixar Animation Studios, Stanford, Sony Research Laboratory, NVIDIA Corporation, Intel Media Research Laboratory
- 2002 and earlier** Interval Research Corporation, UCLA, UC Davis, UCSD, Sacramento State University, San Francisco State University

TUTORIALS AND COURSES

- Panel Moderator, “Industry Panel on Future Directions in Parallel Processing”, IS&T/SPIE Electronic Imaging: Parallel Processing for Imaging Applications, (Jan. 2011)
- IEEE/ACM Supercomputing, co-instructor, full-day tutorials
- “High Performance Computing with CUDA” (November 2009)
 - “High Performance Computing on GPUs with CUDA” (November 2007)
 - “General Purpose Computation on Graphics Hardware” (November 2006)
- ACM SIGGRAPH, co-instructor, half-day course, “Beyond Programmable Shading: Fundamentals” (Aug. 2008)
- International Ph.D. School in Algorithms for Advanced Processor Architectures, The IT University of Copenhagen, Denmark, Instructor (June 2008)
- Panel Moderator, “GPUs vs. Multicore CPUs: On a Converging Course or Fundamentally Different?”, Graphics Hardware 2008 (June 2008)
- Thirteenth Annual Conference on Architectural Support for Programming Languages and Operating Systems (ACM ASPLOS '08), co-instructor, half-day course, “CUDA: A Heterogeneous Parallel Programming Model for Manycore Computing” (March 2008)
- ACM SIGGRAPH, co-instructor, full-day course, “General Purpose Computation on Graphics Hardware” (Aug. 2007)
- IEEE Visualization, co-instructor, full-day tutorial, “General Purpose Computation on Graphics Hardware” (Oct. 2004, Oct. 2005)

PROFESSIONAL REVIEWING

- Conference** ACM SIGGRAPH, High Performance Graphics, Graphics Hardware, Eurographics, ACM SIGGRAPH Asia, ACM Symposium on Interactive 3D Graphics and Games, IEEE/ACM Supercomputing, International European Conference on Parallel and Distributed Computing, Eurographics Symposium on Rendering, International Symposium on Computer Architecture, ACM Symposium on Parallelism in Algorithms and Architectures, IEEE/ACM International Symposium on Microarchitecture, SIBGRAPI, IEEE International Symposium on Information Theory, Symposium on Application Accelerators in High Performance Computing, International Workshop on Characterizing Applications for Heterogeneous Exascale Systems, European Symposium on Algorithms, African Conference on Computational Mechanics
- Journal** IEEE Transactions on Visualization and Computer Graphics, IEEE Computer Graphics and Applications, Computer Graphics Forum, Proceedings of the IEEE, IEEE Transactions on Computers, ACM Transactions on Parallel Computing, IEEE Transactions on Parallel and Distributed Systems, IEEE Computer Architecture Letters, Journal of Computer

Graphics Techniques, IEEE Access, IEEE Transactions on Circuits and Systems for Video Technology, IEEE Micro, Algorithmica, IEEE Transactions on Knowledge and Data Engineering, SpringerPlus, IEEE Transactions on Mobile Computing, IEEE Transactions on Very Large Scale Integration Systems, IEEE/ACM Transactions on Networking, IEEE Software, ACM Transactions on Architecture and Code Optimization, journal of graphics, gpu, and game tools, ACM Transactions on Mathematical Software, The Computer Journal, Journal of Computer Science and Technology, Journal of Low Power Electronics, International Journal of High Performance Computing Applications, IET Image Processing, Computing in Science & Engineering, Journal of Computer Science and Technology, Computers and Electronics in Agriculture

Books Addison-Wesley, Elsevier, GPU Computing Gems

Government and Nonprofit National Science Foundation (panelist and ad hoc), UC MICRO, Department of Energy SBIR/STTR Program, Department of Energy Early Career Research Program, Natural Sciences and Engineering Research Council of Canada, National Science Centre of Poland, Maryland Industrial Partnerships Program, US Civilian Research and Development Foundation, US-Israel Binational Science Foundation, France-Berkeley Fund

**UNIVERSITY
TEACHING**

- EEC 170, Introduction to Computer Architecture (undergraduate)
— Fall 2025, Winter 2024, Winter 2023, Winter 2022, Winter 2021, Fall 2019, Fall 2010, Fall 2007, Winter 2006, Fall 2003
- EEC 171, Parallel Architecture (undergraduate)
— Spring 2019, Spring 2018, Spring 2017, Spring 2014, Spring 2013, Spring 2012, Spring 2011, Spring 2010, Spring 2009, Spring 2008, Spring 2007
- EEC 172, Embedded Systems (undergraduate)
— Fall 2005, Winter 2005, Fall 2004
- EEC 196, Issues in Engineering Design (undergraduate)
— Fall 2018
- EEC 180A, Digital Systems (undergraduate)
— Winter 2009
- EEC 277, Graphics Architecture (graduate)
— Winter 2017, Winter 2015, Winter 2014, Winter 2013, Winter 2012, Winter 2011, Winter 2010, Winter 2009, Winter 2008, Winter 2007, Winter 2006, Winter 2005, Winter 2004, Spring 2003
- EEC 279 (formerly 289Q), Modern Parallel Computing (graduate)
— Fall 2022, Winter 2020, Winter 2018, Winter 2016
- EEC 289Q, Performance Engineering of Software Systems
— Fall 2023, Fall 2021

**PH.D. GRADUATES
SUPERVISED**

- Yuxin Chen (Ph.D., computer science; completed October 2024). Dissertation title: “PGAS-based Fine-Grained Asynchronous Execution on GPUs”. First employment: Luminary Cloud.
- Ahmed Mahmoud (Ph.D., electrical and computer engineering; completed June 2024). Dissertation title: “Unstructured Geometric Data Processing on the GPU: Data Structures & Programming Models”. First employment: Autodesk Research. Current position: postdoc, MIT.
- Zhongyi Lin (Ph.D., electrical and computer engineering; completed September 2023). Dissertation title: “Performance Modeling and Optimization for Machine Learning Workloads”. First employment: AMD. Current employment: Apple.
- Afton Geil (Ph.D., electrical and computer engineering; completed December 2022). Dissertation title: “Managing GPU Memory Size Limitations”. NSF Graduate Research Fellowship. Current employment: Lawrence Berkeley National Laboratory.
- Muhammad Osama (Ph.D., electrical and computer engineering, completed December 2022). Dissertation title: “Load Balancing on the GPU”. First employment: AMD Research.

Jason Mak (Ph.D., computer science, completed September 2022). Dissertation title: “Large-Scale 3D Reconstruction on the GPU”.

Muhammad Awad (Ph.D., electrical and computer engineering, completed June 2022). Dissertation title: “Fully Concurrent GPU Data Structures”. First employment: AMD Research.

Kerry A. Seitz, Jr. (Ph.D., computer science, completed March 2022). Dissertation title: “Toward Unified Shader Programming”. NSF Graduate Research Fellowship. First employment: Weta Digital / Unity Technologies.

Leyuan Wang (Ph.D., computer science, completed March 2020). Dissertation title: “Parallel Algorithms on Graph matching”. First employment: Amazon Web Services. Current employment: Bytedance.

Carl Yang (Ph.D., electrical and computer engineering, completed June 2019). Dissertation title: “High-Performance Linear Algebra-based Graph Framework on the GPU”. First employment: Waymo.

Yuechao Pan (Ph.D., electrical and computer engineering, completed March 2019). Dissertation title: “Multi-GPU Graph Processing”. First employment: Google Brain.

Saman Ashkiani (Ph.D., electrical and computer engineering, completed December 2017). Dissertation title: “Parallel Algorithms and Dynamic Data Structures on the Graphics Processing Unit: a warp-centric approach”. NVIDIA Graduate Fellowship, 2016–17. First employment: OmniSci (MapD), then Google. Current employment: NVIDIA.

Yangzihao Wang (Ph.D., computer science, completed December 2016). Dissertation title: “Gunrock: A Programming Model and Implementation for Graph Analytics on Graphics Processing Units”. NVIDIA Graduate Fellowship Finalist, 2013–14. First employment: Google Brain, then Tencent, then WeChat. Current employment: Sea AI Lab, Singapore.

Andrew Davidson (Ph.D., electrical and computer engineering, completed January 2015). Dissertation title: “Divide-And-Conquer Methods for Data Parallel Machines”. Richard Dorf Graduate Student Award (outstanding ECE graduate student), 2013. First employment: YouTube (Google).

Pinar Muyan-Özçelik (Ph.D., computer science, completed June 2014). Dissertation title: “Running Real-time Tasks on Embedded Systems using GPU Computing”. NVIDIA Graduate Fellowship, 2011–12; Google Anita Borg Memorial Scholarship Finalist, 2009. First employment: assistant professor, computer science, Sacramento State University.

Anjul Patney (Ph.D., electrical and computer engineering, completed December 2013). Dissertation title: “Programmable Graphics Pipelines”. 2014 Zuhair A. Munir Award, honorable mention (2nd place), best doctoral dissertation, UC Davis College of Engineering. NVIDIA Graduate Fellowships, 2009–10, 2010–11. Intel Ph.D. Fellowship, 2011–12. First employment: NVIDIA Research, then Facebook Reality Labs. Current employment: NVIDIA Research.

Stanley Tzeng (Ph.D., computer science, completed September 2013). Dissertation title: “Scheduling on Manycore and Heterogeneous Graphics Processors”. First employment: NVIDIA. Current employment: Google.

Kshitij Gupta (Ph.D., electrical and computer engineering, completed September 2012). Dissertation title: “GPU-based Parallel Application Design for Emerging Mobile Devices”. First employment: Aptina. Current employment: BabbleLabs.

Jeff Stuart (Ph.D., computer science, completed June 2012). Dissertation title: “High-Performance Computing Software Libraries on Modern GPUs”. First employment: Google Zürich.

Yao Zhang (Ph.D., electrical and computer engineering, completed January 2012). Dissertation title: “Performance Modeling for GPU Architectures”. First employment: postdoc, University of Chicago, then postdoc, Argonne National Laboratory, then Google Brain, then Ant Financial. Current employment: Microsoft.

Shubhabrata Sengupta (Ph.D., computer science, completed December 2010). Dissertation title: “Efficient Primitives and Algorithms for Many-core architectures”. 2008 Best Graduate Researcher Award, UC Davis Computer Science Department. NVIDIA Graduate Fellowships,

2007–8, 2008–9. First employment: Intel Research, then Shazam, then Baidu Research. Current employment: Facebook AI Research Lab.

Eric Lengyel (Ph.D., computer science, completed March 2010). Dissertation title: “Voxel-Based Terrain for Real-Time Virtual Simulations”. First employment: Terathon Software LLC.

Leo Szumel (Ph.D., electrical and computer engineering, completed September 2008). Dissertation title: “On The Programming of Emergent Sensor Network Systems”. First employment: Sentilla. Current employment: Virtual Instruments.

Aaron Lefohn (Ph.D., computer science, completed June 2006). Dissertation title: “Glift: Generic Data Structures for Graphics Hardware”. National Science Foundation Graduate Research Fellowship. 2006 Best Doctoral Dissertation Award, UC Davis Computer Science Department. First employment: Neoptica, acquired by Intel in Oct. 2007. Current employment: director of real-time rendering research, NVIDIA Research.

M.S. GRADUATES SUPERVISED

Mythreya K (M.S., computer science, completed September 2024). Thesis title: “Using modern C++ to improve CUDA programs”. First employment: AMD.

Changcheng (Eric) Yuan (M.S., computer science, completed September 2024). Thesis title: “Fast Sparse Matrix Reordering on GPU for Cholesky Based Solvers”. Current position: Ph.D. student, Texas A&M University.

Annie Robison (M.S., computer science, completed February 2024). Project title: “Evaluating and Improving Performance in a GPU Graph Framework”. Current employment: Hamilton Company.

Charles Rozhon (M.S., computer science, completed December 2023). Thesis title: “Capturing & Modifying GPU Binaries in User-mode”. Current employment: Obsidian Entertainment.

Marjerie Suresh (M.S., electrical and computer engineering, completed September 2023). Thesis title: “Dynamically Determined Output Space For SpGEMM Using GPU Virtual Memory Management”. Current employment: Samsung.

Jonathan Wapman (M.S., electrical and computer engineering, completed December 2022). Thesis title: “Harmonic CUDA: Asynchronous Programming on GPUs”. Current employment: Waymo.

Weitang Liu (M.S., electrical and computer engineering, completed December 2018). Thesis title: “Sample efficient DQN and Object Localization with Capsnet”. Currently Ph.D. student at UC San Diego.

Matthew Yih (M.S., electrical and computer engineering, completed December 2018). Thesis title: “FPGA versus GPU for Autonomous Vehicle Workload”. First employment: Amazon. Current employment: Google.

Yuduo Wu (M.S., electrical and computer engineering, completed September 2015). Thesis title: “Performance Characterization of High-Level Programming Models for GPU Graph Analytics”. First employment: IBM, then NovuMind Inc. Current employment: Apple.

Ritesh Patel (M.S., electrical and computer engineering, completed June 2012). Thesis title: “Parallel Lossless Data Compression on the GPU”. First employment: Intel, then Velo3D, then Intel. Current employment: Google.

Will Kohut (M.S., electrical and computer engineering, completed June 2011). Thesis title: “Protein All-Atom Energy Evaluation on a GPU”. First employment: NVIDIA. Current employment: Empathic Therapeutics.

Everett Phillips (M.S., mechanical and aerospace engineering, completed March 2010). Co-advised with Dr. Roger Davis, UCD MAE. First employment: NVIDIA.

Vladimir Glavtchev (M.S., electrical and computer engineering, completed December 2009). Thesis title: “EU Speed-limit Sign Detection Using a Graphics Processing Unit (GPU)”. First employment: BMW Group Technology Office, Palo Alto. Subsequent employment: NVIDIA, Fyusion, Marble, Tanium. Current employment: Zoox.

Adam Moerschell (M.S., electrical and computer engineering, completed March 2007). Thesis title: “Distributed Texture Memory in a Multi-GPU Environment”. First employment: ATI, acquired by AMD. Current employment: Apple.

Scott Goering (M.S., electrical and computer engineering, completed March 2005). Thesis title: “Performance of On-Demand Routing Algorithms in Dense, Constrained Networks”. First employment: Intel.

CURRENT STUDENTS

Currently primary advisor for: Mingun Cho (Ph.D., computer science); Daniel Loran (M.S., electrical and computer engineering); Agnieszka Łupińska (Ph.D., computer science); Collin McCarthy (Ph.D., computer science); Toluwanimi Odemuyiwa (Ph.D., electrical and computer engineering); Cameron Shinn (Ph.D., electrical and computer engineering); Radoyeh Shojaei (Ph.D., computer science); Chenfei Yao (M.S., computer science).

REFEREED PUBLICATIONS

Thomas Smith, Raph Levien, and John D. Owens. “Decoupled Fallback: A Portable Single-Pass GPU Scan”. In *Proceedings of the 37th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA ’25, pages 255–268. July 2025.

Ahmed H. Mahmoud, Serban D. Porumbescu, and John D. Owens. “Dynamic Mesh Processing on the GPU (Abstract)”. In *Proceedings of the 3rd Highlights of Parallel Computing Workshop*, HOPC ’25, pages 7–9. July 2025.

Ahmed H. Mahmoud, Serban D. Porumbescu, and John D. Owens. “Dynamic Mesh Processing on the GPU”. *ACM Transactions on Graphics*, 44(4):136:1–19, July 2025.

Zhongyi Lin, Ning Sun, Pallab Bhattacharya, Xizhou Feng, Louis Feng, and John D. Owens. “Towards Universal Performance Modeling for Machine Learning Training on Multi-GPU Platforms”. *IEEE Transactions on Parallel and Distributed Systems*, 36(2):226–238, February 2025.

Yuxin Chen, Aydın Buluç, Katherine Yelick, and John D. Owens. “Accelerating Multi-GPU Embedding Retrieval with PGAS-Style Communication for Deep Learning Recommendation Systems”. In *Parallel Applications Workshop, Alternatives To MPI+X*, PAW-ATM2024, pages 1262–1273. November 2024.

John D. Owens and Bruce Hoppe. “Helping Faculty Teach Software Performance Engineering”. In *Proceedings of the 14th NSF/TCPP Workshop on Parallel and Distributed Computing Education*, EduPar-24, pages 338–341. May 2024.

Scott Rixner, William J. Dally, Ujval J. Kapasi, Peter Mattson, and John D. Owens. “RETROSPECTIVE: Memory Access Scheduling”. In José F. Martínez and Lizy K. John, editors, *ISCA@50 25-Year Retrospective: 1996–2020*. ACM SIGARCH and IEEE TCCA, June 2023.

Afton Geil, Serban D. Porumbescu, and John D. Owens. “Maximum Clique Enumeration on the GPU”. In *Proceedings of the Workshop on Graphs, Architectures, Programming, and Learning*, GrAPL 2023, pages 234–244. May 2023.

Toluwanimi O. Odemuyiwa, Hadi Asghari-Moghaddam, Michael Pellauer, Kartik Hegde, Po-An Tsai, Neal Crago, Aamer Jaleel, John D. Owens, Edgar Solomonik, Joel Emer, and Christopher Fletcher. “Accelerating Sparse Data Orchestration via Dynamic Reflexive Tiling”. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, volume 3 of ASPLOS ’23, pages 18–32. March 2023.

Jonathan D. Wapman, Sean Treichler, Serban D. Porumbescu, and John D. Owens. “Harmonic CUDA: Asynchronous Programming on GPUs”. In *Proceedings of the 14th International Workshop on Programming Models and Applications for Multicores and Manycores*, PMAM ’23, pages 39–49. February 2023.

Muhammad Osama, Duane Merrill, Cris Cecka, Michael Garland, and John D. Owens. “Stream-K: Work-Centric Parallel Decomposition for Dense Matrix-Matrix Multiplication on the GPU”. In *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, PPOPP ’23, pages 429–431. ACM, February 2023.

Muhammad Osama, Serban D. Porumbescu, and John D. Owens. “A Programming Model for GPU Load Balancing”. In *Proceedings of the 28th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP 2023, pages 79–91. February/March 2023.

Muhammad A. Awad, Saman Ashkiani, Serban D. Porumbescu, Martín Farach-Colton, and John D. Owens. “Analyzing and Implementing GPU Hash Tables”. In *SIAM Symposium on Algorithmic Principles of Computer Systems*, APOCS23, pages 33–50. January 2023.

Zhongyi Lin, Louis Feng, Ehsan K. Ardestani, Jaewon Lee, John Lundell, Changkyu Kim, Arun Kejariwal, and John D. Owens. “Building a Performance Model for Deep Learning Recommendation Model Training on GPUs”. In *2022 IEEE 29th International Conference on High Performance Computing, Data, and Analytics*, HiPC 2022, pages 48–58. IEEE, December 2022.

Yuxin Chen, Benjamin Brock, Serban Porumbescu, Aydın Buluç, Katherine Yelick, and John D. Owens. “Scalable Irregular Parallelism with GPU: Getting CPUs Out of the Way”. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, SC ’22, pages 708–723. November 2022.

Muhammad A. Awad, Serban D. Porumbescu, and John D. Owens. “A GPU Multiversion B-Tree”. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, PACT 2022, pages 481–493. October 2022.

Yuxin Chen, Benjamin Brock, Serban Porumbescu, Aydın Buluç, Katherine Yelick, and John D. Owens. “Atos: A Task-Parallel GPU Scheduler for Graph Analytics”. In *Proceedings of the International Conference on Parallel Processing*, ICPP 2022. August/September 2022.

Kerry A. Seitz, Jr., Theresa Foley, Serban D. Porumbescu, and John D. Owens. “Supporting Unified Shader Specialization by Co-opting C++ Features”. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 5(3):25:1–25:17, July 2022.

Muhammad Osama, Serban D. Porumbescu, and John D. Owens. “Essentials of Parallel Graph Analytics”. In *Proceedings of the Workshop on Graphs, Architectures, Programming, and Learning*, GrAPL 2022, pages 314–317. May 2022.

Carl Yang, Aydın Buluç, and John D. Owens. “GraphBLAST: A High-Performance Linear Algebra-based Graph Framework on the GPU”. *ACM Transactions on Mathematical Software*, 48(1):1:1–1:51, February 2022.

Zhongyi Lin, Evangelos Georganas, and John D. Owens. “Towards Flexible and Compiler-friendly Layer Fusion for CNNs on Multi-core CPUs”. In *Euro-Par 2021: Proceedings of the 27th International European Conference on Parallel and Distributed Computing*. September 2021.

Ahmed H. Mahmoud, Serban D. Porumbescu, and John D. Owens. “RXMesh: A GPU Mesh Data Structure”. *ACM Transactions on Graphics*, 40(4):104:1–104:16, August 2021.

Weitang Liu, Xiaoyun Wang, John D. Owens, and Yixuan Li. “Energy-based Out-of-distribution Detection”. In *Advances in Neural Information Processing Systems*, volume 33 of *NeurIPS 2020*. December 2020.

Muhammad A. Awad, Saman Ashkiani, Serban D. Porumbescu, and John D. Owens. “Dynamic Graphs on the GPU”. In *Proceedings of the 34th IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2020, pages 739–748. May 2020.

Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, Ahmed H. Mahmoud, Scott A. Mitchell, John D. Owens, and Ahmad A. Rushdi. “VoroCrust: Voronoi Meshing Without Clipping”. *ACM Transactions on Graphics*, 39(3):23:1–23:16, May 2020.

Kerry A. Seitz, Jr., T. Foley, Serban D. Porumbescu, and John D. Owens. “Staged Metaprogramming for Shader System Development”. *ACM Transactions on Graphics*, 38(6):202:1–202:15, November 2019.

Benjamin Brock, Yuxin Chen, Jiakun Yan, John D. Owens, Aydın Buluç, and Katherine Yelick. “RDMA vs. RPC for Implementing Distributed Data Structures”. In *Proceedings of the IEEE/ACM 9th Workshop on Irregular Applications: Architectures and Algorithms*, IA³ 2019, pages 17–22. November 2019.

Leyuan Wang and John D. Owens. “Fast BFS-Based Triangle Counting on GPUs”. In *Proceedings of the IEEE High Performance Extreme Computing Conference, HPEC '19*. September 2019. 2019 GraphChallenge Finalist.

Xiaoyun Wang, Zhongyi Lin, Carl Yang, and John D. Owens. “Accelerating DNN Inference with GraphBLAS and the GPU”. In *Proceedings of the IEEE High Performance Extreme Computing Conference, HPEC '19*. September 2019. 2019 GraphChallenge Student Innovation Award.

Zhongyi Lin, Matthew Yih, Jeffrey M. Ota, John D. Owens, and Pınar Muyan-Özçelik. “Benchmarking Deep Learning Frameworks and Investigating FPGA Deployment for Traffic Sign Classification and Detection”. *IEEE Transactions on Intelligent Vehicles*, 4(3):385–395, September 2019.

Muhammad Osama, Minh Truong, Carl Yang, Aydın Buluç, and John D. Owens. “Graph Coloring on the GPU”. In *Proceedings of the Workshop on Graphs, Architectures, Programming, and Learning, GrAPL 2019*, pages 231–240. May 2019.

Muhammad A. Awad, Saman Ashkiani, Rob Johnson, Martín Farach-Colton, and John D. Owens. “Engineering a High-Performance GPU B-Tree”. In *Proceedings of the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2019*, pages 145–157. February 2019.

Matthew Yih, Jeffrey M. Ota, John D. Owens, and Pınar Muyan-Özçelik. “FPGA versus GPU for Speed-Limit-Sign Recognition”. In *Proceedings of the 21st IEEE International Conference on Intelligent Transportation Systems, ITSC 2018*, pages 843–850. November 2018.

Carl Yang, Aydın Buluç, and John D. Owens. “Implementing Push-Pull Efficiently in GraphBLAS”. In *Proceedings of the International Conference on Parallel Processing, ICPP 2018*, pages 89:1–89:11. August 2018.

Carl Yang, Aydın Buluç, and John D. Owens. “Design Principles for Sparse Matrix Multiplication on the GPU”. In Marco Aldinucci, Luca Padovani, and Massimo Torquati, editors, *Euro-Par 2018: Proceedings of the 24th International European Conference on Parallel and Distributed Computing*, pages 672–687. August 2018. Distinguished Paper and Best Artifact Award.

John D. Owens. “Technical Perspective: Graphs, Betweenness Centrality, and the GPU”. *Communications of the ACM*, 61(8):84, August 2018.

Zhongyi Lin, Jeffrey M. Ota, John D. Owens, and Pınar Muyan-Özçelik. “Benchmarking Deep Learning Frameworks with FPGA-suitable Models on a Traffic Sign Dataset”. In *Proceedings of the 2018 IEEE Intelligent Vehicles Symposium, IV '18*, pages 1197–1203. June 2018.

Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, Ahmed H. Mahmoud, Scott A. Mitchell, John D. Owens, and Ahmad A. Rushdi. “VoroCrust Illustrated: Theory and Challenges (Multimedia Exposition)”. In Bettina Speckmann and Csaba D. Tóth, editors, *34th International Symposium on Computational Geometry (SoCG 2018)*, volume 99 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 77:1–77:4. Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, June 2018.

Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, Ahmed H. Mahmoud, Scott A. Mitchell, John D. Owens, and Ahmad Rushdi. “Sampling Conditions for Conforming Voronoi Meshing by the VoroCrust Algorithm”. In Bettina Speckmann and Csaba D. Tóth, editors, *34th International Symposium on Computational Geometry (SoCG 2018)*, volume 99 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1:1–1:16. Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, June 2018.

Yuechao Pan, Roger Pearce, and John D. Owens. “Scalable Breadth-First Search on a GPU Cluster”. In *Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium, IPDPS 2018*, pages 1090–1101. May 2018.

Afton Geil, Martin Farach-Colton, and John D. Owens. “Quotient Filters: Approximate Membership Queries on the GPU”. In *Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium, IPDPS 2018*, pages 451–462. May 2018.

- Saman Ashkiani, Shengren Li, Martin Farach-Colton, Nina Amenta, and John D. Owens. “GPU LSM: A Dynamic Dictionary Data Structure for the GPU”. In *Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2018, pages 430–440. May 2018.
- Saman Ashkiani, Martin Farach-Colton, and John D. Owens. “A Dynamic Hash Table for the GPU”. In *Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2018, pages 419–429. May 2018.
- Yangzihao Wang, Yuechao Pan, Andrew Davidson, Yuduo Wu, Carl Yang, Leyuan Wang, Muhammad Osama, Chenshan Yuan, Weitang Liu, Andy T. Riffel, and John D. Owens. “Gunrock: GPU Graph Analytics”. *ACM Transactions on Parallel Computing*, 4(1):3:1–3:49, August 2017.
- Pınar Muyan-Özçelik and John D. Owens. “Methods for Multitasking among Real-time Embedded Compute Tasks Running on the GPU”. *Concurrency and Computation: Practice and Experience*, 29(15):e4118:1–e4118:14, August 2017.
- Saman Ashkiani, Andrew A. Davidson, Ulrich Meyer, and John D. Owens. “GPU Multisplit: an extended study of a parallel algorithm”. *ACM Transactions on Parallel Computing*, 4(1):2:1–2:44, August 2017.
- Ahmed Abdelkader, Ahmed H. Mahmoud, Ahmad A. Rushdi, Scott A. Mitchell, John D. Owens, and Mohamed S. Ebeida. “A Constrained Resampling Strategy for Mesh Improvement”. *Computer Graphics Forum*, 36(5):189–201, July 2017.
- Yangzihao Wang, Sean Baxter, and John D. Owens. “Mini-Gunrock: A Lightweight Graph Analytics Framework on the GPU”. In *Graph Algorithms Building Blocks*, GABB 2017, pages 616–626. May 2017.
- Yuechao Pan, Yangzihao Wang, Yuduo Wu, Carl Yang, and John D. Owens. “Multi-GPU Graph Analytics”. In *Proceedings of the 31st IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2017, pages 479–490. May/June 2017.
- David Luebke and John Owens. “Pixels at Scale: High-Performance Computer Graphics and Vision”. In *Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2016 Symposium*, pages 3–5. The National Academies Press, 2017.
- Jonathan Y. Kemal, Roger L. Davis, and John D. Owens. “Multidisciplinary simulation acceleration using multiple shared memory graphical processing units”. *International Journal of High Performance Computing Applications*, 30(4):486–508, November 2016.
- Ross K. Gegan, Vishal Ahuja, John D. Owens, and Dipak Ghosal. “Real-Time GPU-based Timing Channel Detection using Entropy”. In *Proceedings of the IEEE Conference on Communications and Network Security*, CNS 2016, pages 296–305. October 2016.
- Jeremy Kepner, Peter Aaltonen, David Bader, Aydın Buluç, Franz Franchetti, John Gilbert, Dylan Hutchison, Manoj Kumar, Andrew Lumsdaine, Henning Meyerhenke, Scott McMillan, Jose Moreira, John D. Owens, Carl Yang, Marcin Zalewski, and Timothy Mattson. “Mathematical Foundations of the GraphBLAS”. In *Proceedings of the IEEE High Performance Extreme Computing Conference*. September 2016.
- Leyuan Wang, Sean Baxter, and John D. Owens. “Fast Parallel Skew and Prefix-Doubling Suffix Array Construction on the GPU”. *Concurrency and Computation: Practice & Experience*, 28(12):3466–3484, 25 August 2016.
- Saman Ashkiani, Nina Amenta, and John D. Owens. “Parallel Approaches to the String Matching Problem on the GPU”. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA 2016, pages 275–285. July 2016.
- Mohamed S. Ebeida, Ahmad A. Rushdi, Muhammad A. Awad, Ahmed H. Mahmoud, Dong-Ming Yan, Shawn A. English, John D. Owens, Chandrajit L. Bajaj, and Scott A. Mitchell. “Disk Density Tuning of a Maximal Random Packing”. *Computer Graphics Forum*, 35(5):259–269, June 2016.
- Leyuan Wang, Yangzihao Wang, Carl Yang, and John D. Owens. “A Comparative Study on Exact

- Triangle Counting Algorithms on the GPU”. In *Proceedings of the 1st High Performance Graph Processing Workshop*, HPGP ’16, pages 1–8. May 2016.
- Yangzihao Wang, Andrew Davidson, Yuechao Pan, Yuduo Wu, Andy Riffel, and John D. Owens. “Gunrock: A High-Performance Graph Processing Library on the GPU”. In *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP 2016, pages 11:1–11:12. March 2016. Distinguished Paper.
- Pınar Muyan-Özçelik and John D. Owens. “Multitasking Real-time Embedded GPU Computing Tasks”. In *Proceedings of the 7th International Workshop on Programming Models and Applications for Multicores and Manycores*, PMAM 2016, pages 78–87. March 2016.
- Saman Ashkiani, Andrew A. Davidson, Ulrich Meyer, and John D. Owens. “GPU Multisplit”. In *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP 2016, pages 12:1–12:13. March 2016.
- Yuduo Wu, Yangzihao Wang, Yuechao Pan, Carl Yang, and John D. Owens. “Performance Characterization of High-Level Programming Models for GPU Graph Analytics”. In *IEEE International Symposium on Workload Characterization*, IISWC-2015, pages 66–75. October 2015. Best Paper finalist.
- Mikhail M. Shashkov, Jason Mak, Shawn Recker, Connie Nguyen, John Owens, and Kenneth I. Joy. “Efficient Dense Reconstruction Using Geometry and Image Consistency Constraints”. In *Proceedings of the IEEE Applied Imagery Pattern Recognition Workshop*, AIPR 2015. October 2015.
- Leyuan Wang, Sean Baxter, and John D. Owens. “Fast Parallel Suffix Array on the GPU”. In *Euro-Par 2015: Proceedings of the 21st International European Conference on Parallel and Distributed Computing*, volume 9233 of *Lecture Notes in Computer Science*, pages 573–587. Springer, August 2015. Distinguished Paper.
- Anjul Patney, Stanley Tzeng, Kerry A. Seitz, Jr., and John D. Owens. “Piko: A Framework for Authoring Programmable Graphics Pipelines”. *ACM Transactions on Graphics*, 34(4):147:1–147:13, August 2015.
- Carl Yang, Yangzihao Wang, and John D. Owens. “Fast Sparse Matrix and Sparse Vector Multiplication Algorithm on the GPU”. In *Graph Algorithms Building Blocks*, GABB 2015, pages 841–847. May 2015.
- Thomas Weber, Michael Wimmer, and John D. Owens. “Parallel Reyes-style Adaptive Subdivision with Bounded Memory Usage”. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, i3D 2015, pages 39–45. February/March 2015.
- Jonathan Kemal, Roger L. Davis, and John D. Owens. “Multidisciplinary Simulation Acceleration using Multiple Shared-Memory Graphical Processing Units”. In *AIAA Infotech @ Aerospace*, AIAA Science and Technology Forum. January 2015.
- Jason Mak, Mauricio Hess-Flores, Shawn Recker, John D. Owens, and Kenneth I. Joy. “A Comparative Study of Recent GPU-Accelerated Multi-View Sequential Reconstruction Triangulation Methods for Large-Scale Scenes”. In C. V. Jawahar and Shiguang Shan, editors, *Big Data in 3D Computer Vision (Computer Vision—ACCV 2014 Workshops)*, volume 9008 of *Lecture Notes in Computer Science*, pages 254–269. Springer International Publishing, November 2014.
- Mohamed Ebeida, Scott Mitchell, Anjul Patney, Andrew Davidson, Stanley Tzeng, Muhammad Awad, Ahmed Mahmoud, and John 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*, pages 221–238. Springer, November 2014.
- Afton Geil, Yangzihao Wang, and John D. Owens. “WTF, GPU! Computing Twitter’s Who-To-Follow on the GPU”. In *Proceedings of the Second ACM Conference on Online Social Networks*, COSN ’14, pages 63–68. October 2014.
- Andrew Davidson, Sean Baxter, Michael Garland, and John D. Owens. “Work-Efficient Parallel

GPU Methods for Single-Source Shortest Paths”. In *Proceedings of the 28th IEEE International Parallel and Distributed Processing Symposium, IPDPS 2014*, pages 349–359. May 2014.

Jason Mak, Mauricio Hess-Flores, Shawn Recker, John D. Owens, and Kenneth I. Joy. “GPU-Accelerated and Efficient Multi-View Triangulation for Scene Reconstruction”. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision, WACV ’14*, pages 61–68. March 2014.

Mohamed S. Ebeida, Anjul Patney, Scott A. Mitchell, Keith R. Dalbey, Andrew A. Davidson, and John D. Owens. “ k -d Darts: Sampling by k -Dimensional Flat Searches”. *ACM Transactions on Graphics*, 33(1):3:1–3:16, January 2014.

Mohamed S. Ebeida, Ahmed H. Mahmoud, Muhammad A. Awad, Mohammed A. Mohammed, Scott A. Mitchell, Alex Rand, and John D. Owens. “Sifted Disks”. *Computer Graphics Forum*, 32(2):509–518, May 2013.

Stanley Tzeng, Brandon Lloyd, and John D. Owens. “A GPU Task-Parallel Model with Dependency Resolution”. *IEEE Computer*, 45(8):34–41, August 2012.

Stanley Tzeng, Anjul Patney, Andrew Davidson, Mohamed S. Ebeida, Scott A. Mitchell, and John D. Owens. “High-Quality Parallel Depth-of-Field Using Line Samples”. In *Proceedings of High Performance Graphics, HPG ’12*, pages 23–31. June 2012.

Shengren Li, Lance Simons, Jagadeesh Bhaskar Pakaravoor, Fatemeh Abbasinejad, John D. Owens, and Nina Amenta. “kANN on the GPU with Shifted Sorting”. In *Proceedings of High Performance Graphics, HPG ’12*, pages 39–47. June 2012.

Ritesh A. Patel, Yao Zhang, Jason Mak, and John D. Owens. “Parallel Lossless Data Compression on the GPU”. In *Proceedings of Innovative Parallel Computing, InPar ’12*. May 2012.

Kshitij Gupta, Jeff Stuart, and John D. Owens. “A Study of Persistent Threads Style GPU Programming for GPGPU Workloads”. In *Proceedings of Innovative Parallel Computing, InPar ’12*. May 2012.

Mohamed S. Ebeida, Scott A. Mitchell, Anjul Patney, Andrew A. Davidson, and John D. Owens. “A Simple Algorithm for Maximal Poisson-Disk Sampling in High Dimensions”. *Computer Graphics Forum*, 31(2):785–794, May 2012.

Andrew Davidson, David Tarjan, Michael Garland, and John D. Owens. “Efficient Parallel Merge Sort for Fixed and Variable Length Keys”. In *Proceedings of Innovative Parallel Computing, InPar ’12*. May 2012.

Andrew Davidson and John Owens. “Toward Techniques for Auto-tuning GPU Algorithms”. In Kristján Jónasson, editor, *Applied Parallel and Scientific Computing*, volume 7134 of *Lecture Notes in Computer Science*, pages 110–119. Springer Berlin / Heidelberg, February 2012.

Yao Zhang, John Ludd Recker, Robert Ulichney, Ingeborg Tastl, and John D. Owens. “Plane-dependent Error Diffusion on a GPU”. In *Proceedings of SPIE: IS&T/SPIE Electronic Imaging 2012 / Parallel Processing for Imaging Applications II*, volume 8295B, pages 8295B–59:1–10. January 2012.

Mohamed S. Ebeida, Anjul Patney, John D. Owens, and Eric Mestreau. “Isotropic conforming refinement of quadrilateral and hexahedral meshes using two-refinement templates”. *International Journal for Numerical Methods in Engineering*, 88(10):974–985, 9 December 2011.

Kshitij Gupta and John D. Owens. “Compute & Memory Optimizations for High-Quality Speech Recognition on Low-End GPU Processors”. In *Proceedings of the International Conference on High Performance Computing, HiPC 2011*. December 2011.

Yao Zhang, Jonathan Cohen, Andrew A. Davidson, and John D. Owens. “A Hybrid Method for Solving Tridiagonal Systems on the GPU”. In Wen-mei W. Hwu, editor, *GPU Computing Gems*, volume 2, chapter 11, pages 117–132. Morgan Kaufmann, October 2011.

Jeff A. Stuart, Pavan Balaji, and John D. Owens. “Extending MPI to Accelerators”. In *Proceedings of the First Workshop on Architectures and Systems for Big Data, ASBD 2011*, pages 19–23. October 2011.

Mark Silberstein, Assaf Schuster, and John D. Owens. "Applying Software-Managed Caching and CPU/GPU Task Scheduling for Accelerating Dynamic Workloads". In Wen-mei W. Hwu, editor, *GPU Computing Gems*, volume 2, chapter 36, pages 501–517. Morgan Kaufmann, October 2011.

Mohamed S. Ebeida, Scott A. Mitchell, Andrew A. Davidson, Anjul Patney, Patrick M. Knupp, and John D. Owens. "Efficient and Good Delaunay Meshes From Random Points". In *Proceedings of the SIAM Conference on Geometric and Physical Modeling*, GD/SPM11, pages 1506–1515. October 2011.

Dan A. Alcantara, Vasily Volkov, Shubhabrata Sengupta, Michael Mitzenmacher, John D. Owens, and Nina Amenta. "Building an Efficient Hash Table on the GPU". In Wen-mei W. Hwu, editor, *GPU Computing Gems*, volume 2, chapter 4, pages 39–53. Morgan Kaufmann, October 2011.

Everett H. Phillips, Yao Zhang, Roger L. Davis, and John D. Owens. "Acceleration of 2-D Compressible Flow Solvers with Graphics Processing Unit Clusters". *Journal of Aerospace Computing, Information, and Communication*, 8(8):237–249, August 2011.

John Jenkins, Isha Arkatkar, John D. Owens, Alok Choudhary, and Nagiza F. Samatova. "Lessons Learned from Exploring the Backtracking Paradigm on the GPU". In *Euro-Par 2011: Proceedings of the 17th International European Conference on Parallel and Distributed Computing*, volume 6853 of *Lecture Notes in Computer Science*, pages 425–437. Springer, August/September 2011.

Jeff A. Stuart, Michael Cox, and John D. Owens. "GPU-to-CPU Callbacks". In *Euro-Par 2010 Workshops: Proceedings of the Third Workshop on UnConventional High Performance Computing (UCHPC 2010)*, volume 6586 of *Lecture Notes in Computer Science*, pages 365–372. Springer, July 2011.

Mohamed S. Ebeida, Anjul Patney, Scott A. Mitchell, Andrew Davidson, Patrick M. Knupp, and John D. Owens. "Efficient Maximal Poisson-Disk Sampling". *ACM Transactions on Graphics*, 30(4):49:1–49:12, July 2011.

Christopher P. Stone, Earl P. N. Duque, Yao Zhang, David Car, John D. Owens, and Roger L. Davis. "GPGPU parallel algorithms for structured-grid CFD codes". In *Proceedings of the 20th AIAA Computational Fluid Dynamics Conference*, 2011-3221. June 2011.

Vladimir Glavtchev, Pinar Muyan-Özçelik, Jeffrey M. Ota, and John D. Owens. "Feature-Based Speed Limit Sign Detection Using a Graphics Processing Unit". In *Proceedings of the 2011 IEEE Intelligent Vehicles Symposium, IV '11*, pages 195–200. June 2011.

Jeff A. Stuart and John D. Owens. "Multi-GPU MapReduce on GPU Clusters". In *Proceedings of the 25th IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2011, pages 1068–1079. May 2011.

Andrew Davidson, Yao Zhang, and John D. Owens. "An Auto-tuned Method for Solving Large Tridiagonal Systems on the GPU". In *Proceedings of the 25th IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2011, pages 956–965. May 2011.

Andrew Davidson and John D. Owens. "Register Packing for Cyclic Reduction: A Case Study". In *Proceedings of the Fourth Workshop on General Purpose Processing on Graphics Processing Units*, GPGPU-4, pages 4:1–4:6. March 2011.

Yao Zhang and John D. Owens. "A Quantitative Performance Analysis Model for GPU Architectures". In *Proceedings of the 17th IEEE International Symposium on High-Performance Computer Architecture*, HPCA-17, pages 382–393. February 2011.

Pinar Muyan-Özçelik, Vladimir Glavtchev, Jeffrey M. Ota, and John D. Owens. "Real-Time Speed-Limit-Sign Recognition on an Embedded System Using a GPU". In Wen-mei W. Hwu, editor, *GPU Computing Gems*, volume 1, chapter 32, pages 497–516. Morgan Kaufmann, February 2011.

Yao Zhang, John Ludd Recker, Robert Ulichney, Giordano B. Beretta, Ingeborg Tastl, I-Jong Lin, and John D. Owens. "A Parallel Error Diffusion Implementation on a GPU". In *Proceedings of SPIE: IS&T/SPIE Electronic Imaging 2011 / Parallel Processing for Imaging Applications*, volume 7872, pages 78720K:1–9. January 2011.

Shubhabrata Sengupta, Mark Harris, Michael Garland, and John D. Owens. “Efficient Parallel Scan Algorithms for many-core GPUs”. In Jakub Kurzak, David A. Bader, and Jack Dongarra, editors, *Scientific Computing with Multicore and Accelerators*, Chapman & Hall/CRC Computational Science, chapter 19, pages 413–442. Taylor & Francis, January 2011.

Pınar Muyan-Özçelik, Vladimir Glavtchev, Jeffery M. Ota, and John D. Owens. “A Template-Based Approach for Real-Time Speed-Limit-Sign Recognition on an Embedded System using GPU Computing”. In Michael Goesele, Stefan Roth, Arjan Kuijper, Bernt Schiele, and Konrad Schindler, editors, *DAGM 2010: Proceedings of the 32nd Annual Symposium of the German Association for Pattern Recognition*, volume 6376 of *Lecture Notes in Computer Science*, pages 162–171. Springer, September 2010.

Stanley Tzeng, Anjul Patney, and John D. Owens. “Task Management for Irregular-Parallel Workloads on the GPU”. In *Proceedings of High Performance Graphics*, HPG ’10, pages 29–37. June 2010. 2019 High Performance Graphics Test of Time Award for the most influential paper from HPG’s 2010 predecessor conferences.

Jeff A. Stuart, Cheng-Kai Chen, Kwan-Liu Ma, and John D. Owens. “Multi-GPU Volume Rendering using MapReduce”. In *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing / The First International Workshop on MapReduce and its Applications*, HPDC ’10 / MAPREDUCE ’10, pages 841–848. June 2010.

Everett H. Phillips, Roger L. Davis, and John D. Owens. “Unsteady Turbulent Simulations on a Cluster of Graphics Processors”. In *Proceedings of the 40th AIAA Fluid Dynamics Conference*, AIAA 2010-5036. June 2010.

Anjul Patney, Stanley Tzeng, and John D. Owens. “Fragment-Parallel Composite and Filter”. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)*, 29(4):1251–1258, June 2010.

Andrew Davidson and John D. Owens. “Toward Techniques for Auto-Tuning GPU Algorithms”. In *State of the Art in Scientific and Parallel Computing*, Para 2010. June 2010.

Yao Zhang, Jonathan Cohen, and John D. Owens. “Fast Tridiagonal Solvers on the GPU”. In *Proceedings of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP 2010, pages 127–136. January 2010.

Kshitij Gupta and John D. Owens. “Three-Layer Optimizations for Fast GMM Computations on GPU-like Parallel Processors”. In *Proceedings of the IEEE Workshop on Automatic Speech Recognition & Understanding*, ASRU 2009, pages 146–151. December 2009.

Dan A. Alcantara, Andrei Sharf, Fatemeh Abbasinejad, Shubhabrata Sengupta, Michael Mitzenmacher, John D. Owens, and Nina Amenta. “Real-Time Parallel Hashing on the GPU”. *ACM Transactions on Graphics*, 28(5):154:1–154:9, December 2009.

Anjul Patney, Mohamed S. Ebeida, and John D. Owens. “Parallel View-Dependent Tessellation of Catmull-Clark Subdivision Surfaces”. In *Proceedings of High Performance Graphics*, HPG ’09, pages 99–108. August 2009.

Luke J. Gosink, Kesheng Wu, E. Wes Bethel, John D. Owens, and Kenneth I. Joy. “Data Parallel Bin-Based Indexing for Answering Queries on Multi-Core Architectures”. In *Proceedings of the 21st International Conference on Scientific and Statistical Database Management*, volume 5566 of *Lecture Notes in Computer Science*, pages 110–129. Springer, June 2009.

Jeff A. Stuart and John D. Owens. “Message Passing on Data-Parallel Architectures”. In *Proceedings of the 23rd IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2009. May 2009.

Brian Budge, Tony Bernardin, Jeff A. Stuart, Shubhabrata Sengupta, Kenneth I. Joy, and John D. Owens. “Out-of-core Data Management for Path Tracing on Hybrid Resources”. *Computer Graphics Forum (Proceedings of Eurographics 2009)*, 28(2):385–396, April 2009.

Everett H. Phillips, Yao Zhang, Roger L. Davis, and John D. Owens. “Rapid Aerodynamic Performance Prediction on a Cluster of Graphics Processing Units”. In *Proceedings of the 47th*

- AIAA Aerospace Sciences Meeting, AIAA 2009-565. January 2009.
- Anjul Patney and John D. Owens. “Real-Time Reyes-Style Adaptive Surface Subdivision”. *ACM Transactions on Graphics*, 27(5):143:1–143:8, December 2008.
- Sanjiv S. Samant, Junyi Xia, Pinar Muyan-Özçelik, and John D. Owens. “High performance computing for deformable image registration: Towards a new paradigm in adaptive radiotherapy”. *Medical Physics*, 35(8):3546–3553, August 2008.
- Pinar Muyan-Özçelik, John D. Owens, Junyi Xia, and Sanjiv S. Samant. “Fast Deformable Registration on the GPU: A CUDA Implementation of Demons”. In *Proceedings of the 2008 International Conference on Computational Science and Its Applications (First Technical Session on UnConventional High Performance Computing)*, UCHPC ’08, pages 223–233. July 2008.
- Mark Silberstein, Assaf Schuster, Dan Geiger, Anjul Patney, and John D. Owens. “Efficient Computation of Sum-products on GPUs Through Software-Managed Cache”. In *Proceedings of the 22nd ACM International Conference on Supercomputing*, ICS ’08, pages 309–318. June 2008.
- John D. Owens, Mike Houston, David Luebke, Simon Green, John E. Stone, and James C. Phillips. “GPU Computing”. *Proceedings of the IEEE*, 96(5):879–899, May 2008.
- Adam Moerschell and John D. Owens. “Distributed Texture Memory in a Multi-GPU Environment”. *Computer Graphics Forum*, 27(1):130–151, March 2008.
- Aaron E. Lefohn, Shubhabrata Sengupta, and John D. Owens. “Resolution-Matched Shadow Maps”. *ACM Transactions on Graphics*, 26(4):20:1–20:17, October 2007.
- John D. Owens, William J. Dally, Ron Ho, D. N. Jayasimha, Stephen W. Keckler, and Li-Shiuan Peh. “Research Challenges for On-Chip Interconnection Networks”. *IEEE Micro*, 27(5):96–108, September/October 2007.
- Shubhabrata Sengupta, Mark Harris, Yao Zhang, and John D. Owens. “Scan Primitives for GPU Computing”. In *Proceedings of the 22nd ACM SIGGRAPH/EUROGRAPHICS Symposium on Graphics Hardware*, GH ’07, pages 97–106. August 2007. Best Paper Award. 2017 High Performance Graphics Test of Time Award for the most influential paper from HPG’s 2007–08 predecessor conferences.
- Mark Harris, Shubhabrata Sengupta, and John D. Owens. “Parallel Prefix Sum (Scan) with CUDA”. In Hubert Nguyen, editor, *GPU Gems 3*, chapter 39, pages 851–876. Addison Wesley, August 2007.
- John D. Owens. “Towards Multi-GPU Support for Visualization”. *Journal of Physics: Conference Series*, 78:012055 (5pp), June 2007.
- Kwan-Liu Ma, Robert Ross, Jian Huang, Greg Humphreys, Nelson Max, Kenneth Moreland, John D. Owens, and Han-Wei Shen. “Ultra-Scale Visualization: Research and Education”. *Journal of Physics: Conference Series*, 78:012088 (6pp), June 2007.
- John D. Owens, David Luebke, Naga Govindaraju, Mark Harris, Jens Krüger, Aaron E. Lefohn, and Tim Purcell. “A Survey of General-Purpose Computation on Graphics Hardware”. *Computer Graphics Forum*, 26(1):80–113, March 2007.
- John D. Owens. “The Installation and Use of OpenType Fonts in \LaTeX ”. *TUGboat: Communications of the \TeX Users Group*, 27(2):112–118, December 2006.
- Adam Moerschell and John D. Owens. “Distributed Texture Memory in a Multi-GPU Environment”. In *Proceedings of the 21st ACM SIGGRAPH/EUROGRAPHICS Symposium on Graphics Hardware*, GH ’06, pages 31–38. September 2006.
- Leo Szumel and John D. Owens. “The Virtual Pheromone Communication Primitive”. In Phillip B. Gibbons, Tarek Abdelzaher, James Aspnes, and Ramesh Rao, editors, *Proceedings of the Second IEEE International Conference on Distributed Computing in Sensor Systems*, volume 4026 of *Lecture Notes in Computer Science*, pages 135–149. Springer, June 2006.
- Shubhabrata Sengupta, Aaron E. Lefohn, and John D. Owens. “A Work-Efficient Step-Efficient Prefix Sum Algorithm”. In *Proceedings of the 2006 Workshop on Edge Computing Using New Commodity Architectures*, pages D–26–27. May 2006.

Aaron E. Lefohn, Shubhabrata Sengupta, Joe Kniss, Robert Strzodka, and John D. Owens. "Glift: Generic Data Structures for the GPU". In *Proceedings of the 2006 Workshop on Edge Computing Using New Commodity Architectures*, pages D-15-16. May 2006.

Sung W. Park, Lars Linsen, Oliver Kreylos, John D. Owens, and Bernd Hamann. "Discrete Sibson Interpolation". *IEEE Transactions on Visualization and Computer Graphics*, 12(2):243-253, March/April 2006.

Aaron E. Lefohn, Joe Kniss, Robert Strzodka, Shubhabrata Sengupta, and John D. Owens. "Glift: Generic, Efficient, Random-Access GPU Data Structures". *ACM Transactions on Graphics*, 25(1):60-99, January 2006.

Sung Park, Lars Linsen, Oliver Kreylos, John D. Owens, and Bernd Hamann. "A Framework for Real-Time Volume Visualization of Streaming Scattered Data". In *Proceedings of the Tenth International Fall Workshop on Vision, Modeling, and Visualization*, VMV 2005, pages 225-232. November 2005.

John D. Owens, David Luebke, Naga Govindaraju, Mark Harris, Jens Krüger, Aaron E. Lefohn, and Tim Purcell. "A Survey of General-Purpose Computation on Graphics Hardware". In *Eurographics 2005, State of the Art Reports*, pages 21-51. August 2005.

Aaron Lefohn, Shubhabrata Sengupta, Joe Kniss, Robert Strzodka, and John D. Owens. "Dynamic Adaptive Shadow Maps on Graphics Hardware". In *Technical Sketches Program, ACM SIGGRAPH*. August 2005.

Joe Kniss, Aaron Lefohn, Shubhabrata Sengupta, Robert Strzodka, and John D. Owens. "Octree Textures on Graphics Hardware". In *Technical Sketches Program, ACM SIGGRAPH*. August 2005.

Leo Szumel, Jason LeBrun, and John D. Owens. "Towards a Mobile Agent Framework for Sensor Networks". In *Proceedings of the Second IEEE Workshop on Embedded Networked Sensors*, EmNetS-II, pages 79-87. May 2005.

John Owens. "Streaming Architectures and Technology Trends". In Matt Pharr, editor, *GPU Gems 2*, chapter 29, pages 457-470. Addison Wesley, March 2005.

Aaron Lefohn, Joe Kniss, and John Owens. "Implementing Efficient Parallel Data Structures on GPUs". In Matt Pharr, editor, *GPU Gems 2*, chapter 33, pages 521-545. Addison Wesley, March 2005.

Andrew T. Riffel, Aaron E. Lefohn, Kiril Vidimce, Mark Leone, and John D. Owens. "Mio: Fast Multipass Partitioning via Priority-Based Instruction Scheduling". In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware*, GH '04, pages 35-44. August 2004.

Ujval J. Kapasi, Scott Rixner, William J. Dally, Brucek Khailany, Jung Ho Ahn, Peter Mattson, and John D. Owens. "Programmable Stream Processors". *IEEE Computer*, 36(8):54-62, August 2003.

Brucek Khailany, William J. Dally, Scott Rixner, Ujval J. Kapasi, John D. Owens, and Brian Towles. "Exploring the VLSI Scalability of Stream Processors". In *Proceedings of the Ninth Annual International Symposium on High-Performance Computer Architecture*, HPCA-9, pages 153-164. February 2003.

Ben Serebrin, John D. Owens, Brucek Khailany, Peter Mattson, Ujval J. Kapasi, Chen H. Chen, Jinyung Namkoong, Stephen P. Crago, Scott Rixner, and William J. Dally. "A Stream Processor Development Platform". In *Proceedings of the IEEE International Conference on Computer Design*, ICCD 2002, pages 303-308. Freiburg, Germany, September 2002.

John D. Owens, Ujval J. Kapasi, Peter Mattson, Brian Towles, Ben Serebrin, Scott Rixner, and William J. Dally. "Media Processing Applications on the Imagine Stream Processor". In *Proceedings of the IEEE International Conference on Computer Design*, ICCD 2002, pages 295-302. Freiburg, Germany, September 2002.

John D. Owens, Brucek Khailany, Brian Towles, and William J. Dally. "Comparing Reyes and OpenGL on a Stream Architecture". In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware*, GH '02, pages 47-56. September 2002.

Ujval J. Kapasi, William J. Dally, Brucek Khailany, John D. Owens, and Scott Rixner. “The Imagine Stream Processor”. In *Proceedings of the IEEE International Conference on Computer Design, ICCD 2002*, pages 282–288. Freiburg, Germany, September 2002.

Ujval J. Kapasi, Peter Mattson, William J. Dally, John D. Owens, and Brian Towles. “Stream Scheduling”. In *Proceedings of the 3rd Workshop on Media and Streaming Processors*, pages 101–106. Austin, TX, 2 December 2001.

Brucek Khailany, William J. Dally, Ujval J. Kapasi, Peter Mattson, Jinyung Namkoong, John D. Owens, Brian Towles, Andrew Chang, and Scott Rixner. “Imagine: Media Processing with Streams”. *IEEE Micro*, 21(2):35–46, March/April 2001.

Ujval J. Kapasi, William J. Dally, Scott Rixner, Peter R. Mattson, John D. Owens, and Brucek Khailany. “Efficient Conditional Operations for Data-parallel Architectures”. In *Proceedings of the 33rd Annual ACM/IEEE International Symposium on Microarchitecture, MICRO-33*, pages 159–170. December 2000.

Peter Mattson, William J. Dally, Scott Rixner, Ujval J. Kapasi, and John D. Owens. “Communication Scheduling”. In *Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS-IX*, pages 82–92. November 2000.

John D. Owens, William J. Dally, Ujval J. Kapasi, Scott Rixner, Peter Mattson, and Ben Mowery. “Polygon Rendering on a Stream Architecture”. In *Proceedings of the ACM SIGGRAPH/Eurographics Workshop on Graphics Hardware, HWWS ’00*, pages 23–32. August 2000.

Brucek Khailany, William J. Dally, Scott Rixner, Ujval J. Kapasi, Peter Mattson, Jin Namkoong, John D. Owens, and Brian Towles. “Imagine: Signal and Image Processing Using Streams”. In *Hotchips 12*. August 2000.

Scott Rixner, William J. Dally, Ujval J. Kapasi, Peter Mattson, and John D. Owens. “Memory Access Scheduling”. In *Proceedings of the 27th International Symposium on Computer Architecture, ISCA-2000*, pages 128–138. June 2000.

Scott Rixner, William J. Dally, Brucek Khailany, Peter Mattson, Ujval Kapasi, and John D. Owens. “Register Organization for Media Processing”. In *Proceedings of the Sixth Annual International Symposium on High-Performance Computer Architecture, HPCA-6*, pages 375–386. January 2000.

Scott Rixner, William J. Dally, Ujval J. Kapasi, Brucek Khailany, Abelardo Lopez-Lagunas, Peter Mattson, and John D. Owens. “A Bandwidth-Efficient Architecture for Media Processing”. In *Proceedings of the 31st Annual ACM/IEEE International Symposium on Microarchitecture, MICRO-31*, pages 3–13. December 1998.

OTHER PUBLICATIONS

Jiayi Yuan, Cameron Shinn, Kai Xu, Jingze Cui, George Klimiashvili, Guangxuan Xiao, Perkh Zheng, Bo Li, Yuxin Zhou, Zhouhai Ye, Weijie You, Tian Zheng, Dominic Brown, Pengbo Wang, Richard Cai, Julien Demouth, John D. Owens, Xia Hu, Song Han, Timmy Liu, and Huizi Mao. “BLASST: Dynamic BLocked Attention Sparsity via Softmax Thresholding”. *CoRR*, abs/2512.12087(2512.12087v1), December 2025.

Toluwanimi O. Odemuyiwa, Joel S. Emer, and John D. Owens. “The EDGE Language: Extended General Einsums for Graph Algorithms”. *CoRR*, abs/2404.11591(2404.11591v1), April 2024.

Cameron Shinn, Collin McCarthy, Saurav Muralidharan, Muhammad Osama, and John D. Owens. “The Sparsity Roofline: Understanding the Hardware Limits of Sparse Neural Networks”. *CoRR*, abs/2310.00496(2310.00496v2), September 2023.

Matthew Drescher, Muhammad A. Awad, Serban D. Porumbescu, and John D. Owens. “BOBA: A Parallel Lightweight Graph Reordering Algorithm with Heavyweight Implications”. *CoRR*, abs/2306.10410(2306.10410v2), June 2023.

Muhammad Osama, Duane Merrill, Cris Cecka, Michael Garland, and John D. Owens. “Stream-K: Work-centric Parallel Decomposition for Dense Matrix-Matrix Multiplication on the GPU”. *CoRR*, abs/2301.03598(2301.03598v1), January 2023.

Muhammad A. Awad, Saman Ashkiani, Serban D. Porumbescu, Martín Farach-Colton, and John D. Owens. “Better GPU Hash Tables”. *CoRR*, abs/2108.07232(2108.07232v3), August 2021.

Leyuan Wang and John D. Owens. “Fast Gunrock Subgraph Matching (GSM) on GPUs”. *CoRR*, abs/2003.01527(2003.01527v1), March 2020.

Weitang Liu, Lifeng Wei, James Sharpnack, and John D. Owens. “Unsupervised Object Segmentation with Explicit Localization Module”. *CoRR*, abs/1911.09228(1911.09228v1), November 2019.

Weitang Liu, Emad Barsoum, and John D. Owens. “Object Localization and Motion Transfer learning with Capsules”. *CoRR*, abs/1805.07706(1805.07706v1), May 2018.

Kerry A. Seitz, Jr., Alex Kennedy, Owen Ransom, Bassam A. Younis, and John D. Owens. “A GPU Implementation for Two-Dimensional Shallow Water Modeling”. *CoRR*, abs/1309.1230(1309.1230v1), September 2013.

Stanley Tzeng and John D. Owens. “Finding Convex Hulls Using Quickhull on the GPU”. *CoRR*, abs/1201.2936(1201.2936v1), January 2012.

Jeff A. Stuart and John D. Owens. “Efficient Synchronization Primitives for GPUs”. *CoRR*, abs/1110.4623(1110.4623v1), October 2011.

Luke J. Gosink, Kesheng Wu, E. Wes Bethel, John D. Owens, and Kenneth I. Joy. “Bin-Hash Indexing: A Parallel Method For Fast Query Processing”. Technical Report LBNL-729E, Lawrence Berkeley National Laboratory, 20 August 2008.

Michael Kass, Aaron Lefohn, and John Owens. “Interactive Depth of Field Using Simulated Diffusion on a GPU”. Technical Report #06-01, Pixar Animation Studios, January 2006.

John D. Owens, Shubhabrata Sengupta, and Daniel Horn. “Assessment of Graphic Processing Units (GPUs) for Department of Defense (DoD) Digital Signal Processing (DSP) Applications”. Technical Report ECE-CE-2005-3, Department of Electrical and Computer Engineering, University of California, Davis, October 2005.

John D. Owens. “GPUs tapped for general computing”. *EE Times*, 13 December 2004.

John D. Owens. “On The Scalability of Sensor Network Routing and Compression Algorithms”. Technical Report ECE-CE-2004-1, Computer Engineering Research Laboratory, University of California, Davis, 2004.

Leo Szumel and John D. Owens. “On the Feasibility of the UC Davis Metanet”. Technical Report ECE-CE-2003-2, Computer Engineering Research Laboratory, University of California, Davis, 2003.

John D. Owens. *Computer Graphics on a Stream Architecture*. Ph.D. thesis, Stanford University, November 2002.

Publication links are available at <https://www.ece.ucdavis.edu/~jowens/pubs.html>.