John D Owens

List of Publications by Year
in descending order

115 papers	7,291 citations	279798 23 h-index	161849 54 g-index
117 all docs	117 docs citations	117 times ranked	4724 citing authors

#	Article	IF	Citations
1	A Survey of General-Purpose Computation on Graphics Hardware. Computer Graphics Forum, 2007, 26, 80-113.	3.0	1,370
2	GPU Computing. Proceedings of the IEEE, 2008, 96, 879-899.	21.3	1,266
3	Memory access scheduling. , 2000, , .		524
4	Research Challenges for On-Chip Interconnection Networks. IEEE Micro, 2007, 27, 96-108.	1.8	312
5	Imagine: media processing with streams. IEEE Micro, 2001, 21, 35-46.	1.8	259
6	Gunrock. , 2016, , .		240
7	A quantitative performance analysis model for GPU architectures. , 2011, , .		161
8	Multi-GPU MapReduce on GPU Clusters. , 2011, , .		153
9	A study of Persistent Threads style GPU programming for GPGPU workloads. , 2012, , .		143
10	Mathematical foundations of the GraphBLAS. , 2016, , .		131
11	Work-Efficient Parallel GPU Methods for Single-Source Shortest Paths. , 2014, , .		128
12	Memory access scheduling. Computer Architecture News, 2000, 28, 128-138.	2.5	108
13	Real-time parallel hashing on the GPU. ACM Transactions on Graphics, 2009, 28, 1-9.	7.2	107
14	Fast tridiagonal solvers on the GPU. , 2010, , .		107
15	Glift. ACM Transactions on Graphics, 2006, 25, 60-99.	7.2	105
16	High performance computing for deformable image registration: Towards a new paradigm in adaptive radiotherapy. Medical Physics, 2008, 35, 3546-3553.	3.0	104
17	Gunrock. ACM Transactions on Parallel Computing, 2017, 4, 1-49.	1.4	84
18	A Simple Algorithm for Maximal Poissonâ€Disk Sampling in High Dimensions. Computer Graphics Forum, 2012, 31, 785-794.	3.0	83

#	Article	IF	Citations
19	Gunrock: a high-performance graph processing library on the GPU. , 2015, , .		76
20	Fast tridiagonal solvers on the GPU. ACM SIGPLAN Notices, 2010, 45, 127-136.	0.2	73
21	Efficient maximal poisson-disk sampling. ACM Transactions on Graphics, 2011, 30, 1-12.	7.2	69
22	Discrete Sibson interpolation. IEEE Transactions on Visualization and Computer Graphics, 2006, 12, 243-253.	4.4	66
23	Parallel lossless data compression on the GPU. , 2012, , .		65
24	An Auto-tuned Method for Solving Large Tridiagonal Systems on the GPU. , 2011, , .		64
25	Efficient computation of sum-products on GPUs through software-managed cache. , 2008, , .		63
26	Fast Deformable Registration on the GPU: A CUDA Implementation of Demons. , 2008, , .		61
27	Efficient conditional operations for data-parallel architectures. , 2000, , .		53
28	Design Principles for Sparse Matrix Multiplication on the GPU. Lecture Notes in Computer Science, 2018, , 672-687.	1.3	49
29	Resolution-matched shadow maps. ACM Transactions on Graphics, 2007, 26, 20.	7.2	48
30	Message passing on data-parallel architectures. , 2009, , .		48
31	Efficient parallel merge sort for fixed and variable length keys. , 2012, , .		48
32	Multi-GPU volume rendering using MapReduce. , 2010, , .		47
33	Gunrock. ACM SIGPLAN Notices, 2016, 51, 1-12.	0.2	47
34	A Comparative Study on Exact Triangle Counting Algorithms on the GPU. , 2016, , .		38
35	A Dynamic Hash Table for the GPU. , 2018, , .		38
36	Real-time Reyes-style adaptive surface subdivision. ACM Transactions on Graphics, 2008, 27, 1-8.	7.2	36

#	Article	IF	Citations
37	Out‐of‐core Data Management for Path Tracing on Hybrid Resources. Computer Graphics Forum, 2009, 28, 385-396.	3.0	36
38	Lessons Learned from Exploring the Backtracking Paradigm on the GPU. Lecture Notes in Computer Science, 2011, , 425-437.	1.3	36
39	Polygon rendering on a stream architecture. , 2000, , .		31
40	Multi-GPU Graph Analytics. , 2017, , .		31
41	Register packing for cyclic reduction. , 2011, , .		30
42	Real-time parallel hashing on the GPU. , 2009, , .		29
43	Efficient and good Delaunay meshes from random points. CAD Computer Aided Design, 2011, 43, 1506-1515.	2.7	29
44	Building an Efficient Hash Table on the GPU. , 2012, , 39-53.		29
45	Fragment-Parallel Composite and Filter. Computer Graphics Forum, 2010, 29, 1251-1258.	3.0	27
46	A GPU Task-Parallel Model with Dependency Resolution. Computer, 2012, 45, 34-41.	1.1	27
47	Engineering a high-performance GPU B-Tree. , 2019, , .		25
48	Isotropic conforming refinement of quadrilateral and hexahedral meshes using two‐refinement templates. International Journal for Numerical Methods in Engineering, 2011, 88, 974-985.	2.8	24
49	Implementing Push-Pull Efficiently in GraphBLAS. , 2018, , .		24
50	Parallel view-dependent tessellation of Catmull-Clark subdivision surfaces. , 2009, , .		23
51	VoroCrust. ACM Transactions on Graphics, 2020, 39, 1-16.	7.2	22
52	Benchmarking Deep Learning Frameworks and Investigating FPGA Deployment for Traffic Sign Classification and Detection. IEEE Transactions on Intelligent Vehicles, 2019, 4, 385-395.	12.7	20
53	Performance Characterization of High-Level Programming Models for GPU Graph Analytics. , 2015, , .		19
54	Fast Sparse Matrix and Sparse Vector Multiplication Algorithm on the GPU. , 2015, , .		19

#	Article	IF	Citations
55	GraphBLAST: A High-Performance Linear Algebra-based Graph Framework on the GPU. ACM Transactions on Mathematical Software, 2022, 48, 1-51.	2.9	18
56	Dynamic adaptive shadow maps on graphics hardware. , 2005, , .		17
57	Distributed Texture Memory in a Multiâ€GPU Environment. Computer Graphics Forum, 2008, 27, 130-151.	3.0	15
58	Communication scheduling. ACM SIGPLAN Notices, 2000, 35, 82-92.	0.2	14
59	Dynamic Graphs on the GPU. , 2020, , .		14
60	Extending MPI to accelerators. , 2011, , .		14
61	Mio. Graphics Hardware, 2004, , .	0.0	13
62	Three-layer optimizations for fast GMM computations on GPU-like parallel processors. , 2009, , .		13
63	Feature-based speed limit sign detection using a graphics processing unit. , 2011, , .		13
64	Piko. ACM Transactions on Graphics, 2015, 34, 1-13.	7.2	13
65	Efficient maximal poisson-disk sampling. , 2011, , .		12
66	Fast parallel skew and prefixâ€doubling suffix array construction on the GPU. Concurrency Computation Practice and Experience, 2016, 28, 3466-3484.	2.2	12
67	Technical perspective: Graphs, betweenness centrality, and the GPU. Communications of the ACM, 2018, 61, 84-84.	4.5	12
68	<i>k</i> -d Darts. ACM Transactions on Graphics, 2014, 33, 1-16.	7.2	11
69	GPU multisplit. , 2016, , .		11
70	Fast Parallel Suffix Array on the GPU. Lecture Notes in Computer Science, 2015, , 573-587.	1.3	11
71	Octree textures on graphics hardware. , 2005, , .		10
72	A parallel error diffusion implementation on a GPU. Proceedings of SPIE, 2011, , .	0.8	10

#	ARTICLE	IF	CITATIONS
73	A Constrained Resampling Strategy for Mesh Improvement. Computer Graphics Forum, 2017, 36, 189-201.	3.0	10
74	Distributed texture memory in a multi-GPU environment. , 2006, , .		10
75	A Template-Based Approach for Real-Time Speed-Limit-Sign Recognition on an Embedded System Using GPU Computing. Lecture Notes in Computer Science, 2010, , 162-171.	1.3	9
76	A Hybrid Method for Solving Tridiagonal Systems on the GPU. , 2012, , 117-132.		9
77	GPU Multisplit. ACM Transactions on Parallel Computing, 2017, 4, 1-44.	1.4	9
78	GPU LSM: A Dynamic Dictionary Data Structure for the GPU. , 2018, , .		9
79	Plane-dependent error diffusion on a GPU. Proceedings of SPIE, 2012, , .	0.8	8
80	WTF, GPU! computing twitter's who-to-follow on the GPU. , 2014, , .		8
81	Disk Density Tuning of a Maximal Random Packing. Computer Graphics Forum, 2016, 35, 259-269.	3.0	8
82	RXMesh. ACM Transactions on Graphics, 2021, 40, 1-16.	7.2	8
83	Compute & memory optimizations for high-quality speech recognition on low-end GPU processors. , 2011, , .		7
84	Real-Time Speed-Limit-Sign Recognition on an Embedded System Using a GPU. , 2011, , 497-515.		7
85	Acceleration of 2-D Compressible Flow Solvers with Graphics Processing Unit Clusters. Journal of Aerospace Computing, Information, and Communication, 2011, 8, 237-249.	0.8	7
86	Scalable Breadth-First Search on a GPU Cluster. , 2018, , .		7
87	Accelerating DNN Inference with GraphBLAS and the GPU. , 2019, , .		7
88	Communication scheduling. Computer Architecture News, 2000, 28, 82-92.	2.5	6
89	Ultra-Scale Visualization: Research and Education. Journal of Physics: Conference Series, 2007, 78, 012088.	0.4	6
90	Multitasking Real-time Embedded GPU Computing Tasks. , 2016, , .		6

#	Article	IF	Citations
91	Quotient Filters: Approximate Membership Queries on the GPU. , 2018, , .		6
92	Fast BFS-Based Triangle Counting on GPUs. , 2019, , .		6
93	The Virtual Pheromone Communication Primitive. Lecture Notes in Computer Science, 2006, , 135-149.	1.3	6
94	Sifted Disks. Computer Graphics Forum, 2013, 32, 509-518.	3.0	5
95	Data Parallel Bin-Based Indexing for Answering Queries on Multi-core Architectures. Lecture Notes in Computer Science, 2009, , 110-129.	1.3	5
96	Parallel Approaches to the String Matching Problem on the GPU. , 2016, , .		4
97	FPGA versus GPU for Speed-Limit-Sign Recognition. , 2018, , .		4
98	Benchmarking Deep Learning Frameworks with FPGA-suitable Models on a Traffic Sign Dataset. , 2018, , .		4
99	GPU-to-CPU Callbacks. Lecture Notes in Computer Science, 2011, , 365-372.	1.3	4
100	<title>Media processing using streams</title>. , 1998, , .		3
101	Sampling Conditions for Conforming Voronoi Meshing by the VoroCrust Algorithm. Leibniz International Proceedings in Informatics, LIPIcs, 2018, 99, .	0.0	3
102	Towards multi-GPU support for visualization. Journal of Physics: Conference Series, 2007, 78, 012055.	0.4	2
103	Parallel Reyes-style adaptive subdivision with bounded memory usage. , 2015, , .		2
104	Methods for multitasking among real–time embedded compute tasks running on the GPU. Concurrency Computation Practice and Experience, 2017, 29, e4118.	2.2	2
105	Communication scheduling. Operating Systems Review (ACM), 2000, 34, 82-92.	1.9	1
106	GPU-accelerated and efficient multi-view triangulation for scene reconstruction. , 2014, , .		1
107	GPU multisplit. ACM SIGPLAN Notices, 2016, 51, 1-13.	0.2	1
108	Applying Software-Managed Caching and CPU/GPU Task Scheduling for Accelerating Dynamic Workloads. , 2012, , 501-517.		0

#	Article	IF	Citations
109	Multidisciplinary Simulation Acceleration using Multiple Shared-Memory Graphical Processing Units. , 2015, , .		0
110	Real-time GPU-based timing channel detection using entropy. , 2016, , .		0
111	Multidisciplinary simulation acceleration using multiple shared memory graphical processing units. International Journal of High Performance Computing Applications, 2016, 30, 486-508.	3.7	0
112	Mini-Gunrock: A Lightweight Graph Analytics Framework on the GPU. , 2017, , .		0
113	Towards Flexible and Compiler-Friendly Layer Fusion for CNNs on Multicore CPUs. Lecture Notes in Computer Science, 2021, , 232-248.	1.3	0
114	RXMesh. ACM Transactions on Graphics, 2021, 40, 1-16.	7.2	0
115	A Comparative Study of GPU-Accelerated Multi-view Sequential Reconstruction Triangulation Methods for Large-Scale Scenes. Lecture Notes in Computer Science, 2015, , 254-269.	1.3	0