

Umit V Catalyurek

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/5088779/publications.pdf>

Version: 2024-02-01

179
papers

4,834
citations

136740

32
h-index

138251

58
g-index

189
all docs

189
docs citations

189
times ranked

4399
citing authors

#	ARTICLE	IF	CITATIONS
1	A Block-Based Triangle Counting Algorithm on Heterogeneous Environments. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 444-458.	4.0	2
2	A Unifying Framework to Identify Dense Subgraphs on Streams: Graph Nuclei to Hypergraph Cores. , 2021, , .		6
3	PIGO: A Parallel Graph Input/Output Library. , 2021, , .		1
4	Column-Segmented Sparse Matrix-Matrix Multiplication on Multicore CPUs. , 2021, , .		2
5	StrainHub: a phylogenetic tool to construct pathogen transmission networks. Bioinformatics, 2020, 36, 945-947.	1.8	30
6	Multilevel Algorithms for Acyclic Partitioning of Directed Acyclic Graphs. SIAM Journal of Scientific Computing, 2019, 41, A2117-A2145.	1.3	18
7	Geometric Mapping of Tasks to Processors on Parallel Computers with Mesh or Torus Networks. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 2018-2032.	4.0	3
8	DeepSparse: A Task-Parallel Framework for SparseSolvers on Deep Memory Architectures. , 2019, , .		1
9	Scalable Triangle Counting on Distributed-Memory Systems. , 2019, , .		11
10	Linear Algebra-Based Triangle Counting via Fine-Grained Tasking on Heterogeneous Environments : (Update on Static Graph Challenge). , 2019, , .		6
11	A resource provisioning framework for bioinformatics applications in multi-cloud environments. Future Generation Computer Systems, 2018, 78, 379-391.	4.9	23
12	SiNA: A Scalable Iterative Network Aligner. , 2018, , .		1
13	Graph Manipulations for Fast Centrality Computation. ACM Transactions on Knowledge Discovery From Data, 2017, 11, 1-25.	2.5	20
14	Acyclic Partitioning of Large Directed Acyclic Graphs. , 2017, , .		16
15	SONIC: streaming overlapping community detection. Data Mining and Knowledge Discovery, 2016, 30, 819-847.	2.4	8
16	Incremental k-core decomposition: algorithms and evaluation. VLDB Journal, 2016, 25, 425-447.	2.7	62
17	Tracing Origins of the <i>Salmonella</i> Bareilly Strain Causing a Food-borne Outbreak in the United States. Journal of Infectious Diseases, 2016, 213, 502-508.	1.9	145
18	Multi-Jagged: A Scalable Parallel Spatial Partitioning Algorithm. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 803-817.	4.0	26

#	ARTICLE	IF	CITATIONS
19	Mitigating bias in planning two-colour microarray experiments. <i>International Journal of Data Mining and Bioinformatics</i> , 2015, 13, 31.	0.1	1
20	Allele-specific imbalance mapping at human orthologs of mouse susceptibility to colon cancer (<i>Scc</i>) loci. <i>International Journal of Cancer</i> , 2015, 137, 2323-2331.	2.3	5
21	Querying Co-regulated Genes on Diverse Gene Expression Datasets Via Biclustering. <i>Methods in Molecular Biology</i> , 2015, 1375, 55-74.	0.4	1
22	Diversifying Citation Recommendations. <i>ACM Transactions on Intelligent Systems and Technology</i> , 2015, 5, 1-21.	2.9	22
23	Phylogenetic visualization of the spread of H7 influenza A viruses. <i>Cladistics</i> , 2015, 31, 679-691.	1.5	8
24	Hypergraph partitioning for multiple communication cost metrics: Model and methods. <i>Journal of Parallel and Distributed Computing</i> , 2015, 77, 69-83.	2.7	39
25	Incremental closeness centrality in distributed memory. <i>Parallel Computing</i> , 2015, 47, 3-18.	1.3	12
26	Fast and High Quality Topology-Aware Task Mapping. , 2015, , .		19
27	Regularizing graph centrality computations. <i>Journal of Parallel and Distributed Computing</i> , 2015, 76, 106-119.	2.7	34
28	An Out-of-Core Task-based Middleware for Data-Intensive Scientific Computing. , 2015, , 647-667.		3
29	Syndromic Surveillance of Infectious Diseases meets Molecular Epidemiology in a Workflow and Phylogeographic Application. <i>Studies in Health Technology and Informatics</i> , 2015, 216, 766-70.	0.2	3
30	Exploring the Future of Out-of-Core Computing with Compute-Local Non-Volatile Memory. <i>Scientific Programming</i> , 2014, 22, 125-139.	0.5	1
31	NUSAP1 influences the DNA damage response by controlling BRCA1 protein levels. <i>Cancer Biology and Therapy</i> , 2014, 15, 533-543.	1.5	35
32	Biomedical image analysis on a cooperative cluster of GPUs and multicores. , 2014, , .		10
33	mrSNP: Software to detect SNP effects on microRNA binding. <i>BMC Bioinformatics</i> , 2014, 15, 73.	1.2	46
34	Wide-area overlay networking to manage science DMZ accelerated flows. , 2014, , .		9
35	Exploiting Geometric Partitioning in Task Mapping for Parallel Computers. , 2014, , .		45
36	Massively multithreaded maxflow for image segmentation on the Cray XMT-2. <i>Concurrency Computation Practice and Experience</i> , 2014, 26, 2836-2855.	1.4	1

#	ARTICLE	IF	CITATIONS
37	Hardware/Software Vectorization for Closeness Centrality on Multi-/Many-Core Architectures. , 2014, , .		7
38	Analysis of Partitioning Models and Metrics in Parallel Sparse Matrix-Vector Multiplication. Lecture Notes in Computer Science, 2014, , 174-184.	1.0	3
39	Performance Evaluation of Sparse Matrix Multiplication Kernels on Intel Xeon Phi. Lecture Notes in Computer Science, 2014, , 559-570.	1.0	54
40	A comparative analysis of biclustering algorithms for gene expression data. Briefings in Bioinformatics, 2013, 14, 279-292.	3.2	196
41	Fast recommendation on bibliographic networks with sparse-matrix ordering and partitioning. Social Network Analysis and Mining, 2013, 3, 1097-1111.	1.9	5
42	Benchmarking short sequence mapping tools. BMC Bioinformatics, 2013, 14, 184.	1.2	170
43	Collective Modes in Light Nuclei from First Principles. Physical Review Letters, 2013, 111, 252501.	2.9	103
44	A survey of pipelined workflow scheduling. ACM Computing Surveys, 2013, 45, 1-36.	16.1	32
45	Incremental algorithms for closeness centrality. , 2013, , .		41
46	A Push-Relabel-Based Maximum Cardinality Bipartite Matching Algorithm on GPUs. , 2013, , .		4
47	Hypergraph Sparsification and Its Application to Partitioning. , 2013, , .		5
48	Exploring the future of out-of-core computing with compute-local non-volatile memory. , 2013, , .		15
49	Betweenness centrality on GPUs and heterogeneous architectures. , 2013, , .		49
50	TheAdvisor. , 2013, , .		25
51	Towards a personalized, scalable, and exploratory academic recommendation service. , 2013, , .		11
52	Masher. , 2013, , .		1
53	PRASE. , 2013, , .		0
54	STREAMER: A distributed framework for incremental closeness centrality computation. , 2013, , .		7

#	ARTICLE	IF	CITATIONS
55	Streaming algorithms for k-core decomposition. Proceedings of the VLDB Endowment, 2013, 6, 433-444.	2.1	107
56	No Core CI calculations for light nuclei with chiral 2- and 3-body forces. Journal of Physics: Conference Series, 2013, 454, 012063.	0.3	11
57	Metagenomic Insights into the Carbohydrate-Active Enzymes Carried by the Microorganisms Adhering to Solid Digesta in the Rumen of Cows. PLoS ONE, 2013, 8, e78507.	1.1	123
58	GPU Accelerated Maximum Cardinality Matching Algorithms for Bipartite Graphs. Lecture Notes in Computer Science, 2013, , 850-861.	1.0	9
59	Shattering and Compressing Networks for Betweenness Centrality. , 2013, , .		25
60	Microarray vs. RNA-Seq. , 2012, , .		2
61	An Early Evaluation of the Scalability of Graph Algorithms on the Intel MIC Architecture. , 2012, , .		49
62	On Shared-Memory Parallelization of a Sparse Matrix Scaling Algorithm. , 2012, , .		2
63	Large-scale ab initio configuration interaction calculations for light nuclei. Journal of Physics: Conference Series, 2012, 403, 012019.	0.3	20
64	Partitioning Hypergraphs in Scientific Computing Applications through Vertex Separators on Graphs. SIAM Journal of Scientific Computing, 2012, 34, A970-A992.	1.3	12
65	Improving performance of adaptive component-based dataflow middleware. Parallel Computing, 2012, 38, 289-309.	1.3	14
66	Merging network patterns: a general framework to summarize biomedical network data. Network Modeling Analysis in Health Informatics and Bioinformatics, 2012, 1, 103-116.	1.2	9
67	Multithreaded Clustering for Multi-level Hypergraph Partitioning. , 2012, , .		16
68	Scalable Hybrid Implementation of Graph Coloring Using MPI and OpenMP. , 2012, , .		10
69	An Out-of-Core Eigensolver on SSD-equipped Clusters. , 2012, , .		8
70	An Out-of-Core Dataflow Middleware to Reduce the Cost of Large Scale Iterative Solvers. , 2012, , .		11
71	Fast Recommendation on Bibliographic Networks. , 2012, , .		16
72	Load-balancing spatially located computations using rectangular partitions. Journal of Parallel and Distributed Computing, 2012, 72, 1201-1214.	2.7	20

#	ARTICLE	IF	CITATIONS
73	Graph coloring algorithms for multi-core and massively multithreaded architectures. <i>Parallel Computing</i> , 2012, 38, 576-594.	1.3	60
74	Optimizing dataflow applications on heterogeneous environments. <i>Cluster Computing</i> , 2012, 15, 125-144.	3.5	16
75	Optimizing the stretch of independent tasks on a cluster: From sequential tasks to moldable tasks. <i>Journal of Parallel and Distributed Computing</i> , 2012, 72, 489-503.	2.7	6
76	On Reducing I/O Overheads in Large-Scale Invariant Subspace Projections. <i>Lecture Notes in Computer Science</i> , 2012, , 305-314.	1.0	1
77	The Zoltan and Isorropia Parallel Toolkits for Combinatorial Scientific Computing: Partitioning, Ordering and Coloring. <i>Scientific Programming</i> , 2012, 20, 129-150.	0.5	79
78	Parallel Partitioning, Coloring and Ordering in Scientific Computing. <i>Chapman & Hall/CRC Computational Science</i> , 2012, , 351-371.	0.5	1
79	Partitioning Spatially Located Computations Using Rectangles. , 2011, , .		6
80	Improving graph coloring on distributed-memory parallel computers. , 2011, , .		6
81	Distributed-Memory Parallel Algorithms for Matching and Coloring. , 2011, , .		14
82	Benchmarking Short Sequence Mapping Tools. , 2011, , .		5
83	Acceleration of Two Point Correlation Function Calculation for Pathology Image Segmentation. , 2011, 2011, 174-181.		1
84	Hypergraph Partitioning-Based Fill-Reducing Ordering for Symmetric Matrices. <i>SIAM Journal of Scientific Computing</i> , 2011, 33, 1996-2023.	1.3	14
85	Dynamic Decision and Data-Driven Strategies for the Optimal Management of Subsurface Geo-Systems. <i>Journal of Algorithms and Computational Technology</i> , 2011, 5, 645-665.	0.4	0
86	The Supramap project: linking pathogen genomes with geography to fight emergent infectious diseases. <i>Cladistics</i> , 2011, 27, 61-66.	1.5	20
87	Optimizing latency and throughput of application workflows on clusters. <i>Parallel Computing</i> , 2011, 37, 694-712.	1.3	13
88	Integrated data placement and task assignment for scientific workflows in clouds. , 2011, , .		51
89	KELVIN: A Software Package for Rigorous Measurement of Statistical Evidence in Human Genetics. <i>Human Heredity</i> , 2011, 72, 276-288.	0.4	32
90	Color and texture analysis using emerging parallel architectures. <i>International Journal of High Performance Computing Applications</i> , 2011, 25, 404-427.	2.4	0

#	ARTICLE	IF	CITATIONS
91	Message from the Workshops Chair. , 2010, , .		0
92	Development of a code-agnostic computational infrastructure for the dynamic generation of accident progression event trees. Reliability Engineering and System Safety, 2010, 95, 278-294.	5.1	53
93	A Matrix Partitioning Interface to PaToH in MATLAB. Parallel Computing, 2010, 36, 254-272.	1.3	8
94	Distributed-Memory Parallel Algorithms for Distance-2 Coloring and Related Problems in Derivative Computation. SIAM Journal of Scientific Computing, 2010, 32, 2418-2446.	1.3	20
95	Automatic dataflow application tuning for heterogeneous systems. , 2010, , .		7
96	On the Scalability of Hypergraph Models for Sparse Matrix Partitioning. , 2010, , .		3
97	Comparative analysis of biclustering algorithms. , 2010, , .		34
98	Run-time optimizations for replicated dataflows on heterogeneous environments. , 2010, , .		26
99	On Two-Dimensional Sparse Matrix Partitioning: Models, Methods, and a Recipe. SIAM Journal of Scientific Computing, 2010, 32, 656-683.	1.3	85
100	Exploring parallelism in short sequence mapping using Burrows-Wheeler Transform. , 2010, , .		5
101	An Image Analysis Approach for Detecting Malignant Cells in Digitized H&E-stained Histology Images of Follicular Lymphoma. , 2010, , .		9
102	A Moldable Online Scheduling Algorithm and Its Application to Parallel Short Sequence Mapping. Lecture Notes in Computer Science, 2010, , 93-109.	1.0	5
103	Data Flow Frameworks for Emerging Heterogeneous Architectures and Their Application to Biomedicine. Chapman & Hall/CRC Computational Science, 2010, , 375-392.	0.5	0
104	High-performance signal processing on emerging many-core architectures using cuda. , 2009, , .		8
105	Parallel short sequence mapping for high throughput genome sequencing. , 2009, , .		15
106	Parallel materialization of large ABoxes. , 2009, , 1257-1261.		3
107	A Knowledge-Anchored Integrative Image Search and Retrieval System. Journal of Digital Imaging, 2009, 22, 166-182.	1.6	7
108	A repartitioning hypergraph model for dynamic load balancing. Journal of Parallel and Distributed Computing, 2009, 69, 711-724.	2.7	74

#	ARTICLE	IF	CITATIONS
109	Histopathological Image Analysis Using Model-Based Intermediate Representations and Color Texture: Follicular Lymphoma Grading. <i>Journal of Signal Processing Systems</i> , 2009, 55, 169-183.	1.4	163
110	Computer-aided prognosis of neuroblastoma on whole-slide images: Classification of stromal development. <i>Pattern Recognition</i> , 2009, 42, 1093-1103.	5.1	188
111	Compaction of Schedules and a Two-Stage Approach for Duplication-Based DAG Scheduling. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2009, 20, 857-871.	4.0	53
112	A component-based framework for the Cell Broadband Engine. , 2009, , .		8
113	Investigating the use of GPU-accelerated nodes for SAR image formation. , 2009, , .		36
114	Identification of a Breast Cancer Associated Regulatory Network. , 2009, , .		1
115	Coordinating the use of GPU and CPU for improving performance of compute intensive applications. , 2009, , .		59
116	An Integrated Approach to Locality-Conscious Processor Allocation and Scheduling of Mixed-Parallel Applications. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2009, 20, 1158-1172.	4.0	36
117	Computer-aided Prognosis of Neuroblastoma: Detection of mitosis and karyorrhexis cells in digitized histological images. , 2009, 2009, 1433-6.		40
118	Stroma classification for neuroblastoma on graphics processors. <i>International Journal of Data Mining and Bioinformatics</i> , 2009, 3, 280.	0.1	19
119	Advances in parallel partitioning, load balancing and matrix ordering for scientific computing. <i>Journal of Physics: Conference Series</i> , 2009, 180, 012008.	0.3	9
120	A Biclustering Method to Discover Co-regulated Genes Using Diverse Gene Expression Datasets. <i>Lecture Notes in Computer Science</i> , 2009, , 151-163.	1.0	27
121	A framework for scalable greedy coloring on distributed-memory parallel computers. <i>Journal of Parallel and Distributed Computing</i> , 2008, 68, 515-535.	2.7	45
122	Dynamic generation of accident progression event trees. <i>Nuclear Engineering and Design</i> , 2008, 238, 3457-3467.	0.8	83
123	Biomedical image analysis on a cooperative cluster of GPUs and multicores. , 2008, , .		64
124	e-Science, caGrid, and Translational Biomedical Research. <i>Computer</i> , 2008, 41, 58-66.	1.2	15
125	Large-Scale Biomedical Image Analysis in Grid Environments. <i>IEEE Transactions on Information Technology in Biomedicine</i> , 2008, 12, 154-161.	3.6	18
126	Texture classification using nonlinear color quantization: Application to histopathological image analysis. <i>Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing</i> , 2008, , .	1.8	34

#	ARTICLE	IF	CITATIONS
127	A Duplication Based Algorithm for Optimizing Latency Under Throughput Constraints for Streaming Workflows. , 2008, , .		8
128	Computer-aided prognosis of neuroblastoma: classification of stromal development on whole-slide images. Proceedings of SPIE, 2008, , .	0.8	10
129	Computerized microscopic image analysis of follicular lymphoma. Proceedings of SPIE, 2008, , .	0.8	21
130	Using overlays for efficient data transfer over shared wide-area networks. , 2008, , .		25
131	Multi-hop path splitting and multi-pathing optimizations for data transfers over shared wide-area networks using gridFTP. , 2008, , .		5
132	BlueGene/L applications: Parallelism On a Massive Scale. International Journal of High Performance Computing Applications, 2008, 22, 33-51.	2.4	3
133	Translational research design templates, Grid computing, and HPC. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, 2008, 1-15.	1.0	2
134	A dynamic scheduling approach for coordinated wide-area data transfers using GridFTP. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	12
135	Combinatorial algorithms for computational science and engineering. Journal of Physics: Conference Series, 2008, 125, 012071.	0.3	0
136	An Efficient and Reliable Scientific Workflow System. , 2007, , .		6
137	Intelligent Optimization of Parallel and Distributed Applications. , 2007, , .		1
138	Performance vs. accuracy trade-offs for large-scale image analysis applications. , 2007, , .		3
139	A global address space framework for locality aware scheduling of block-sparse computations. , 2007, , .		0
140	Enabling high performance computational science through combinatorial algorithms. Journal of Physics: Conference Series, 2007, 78, 012058.	0.3	0
141	Flexible patient information search and retrieval framework: pilot implementation. , 2007, , .		1
142	Hypergraph-based Dynamic Load Balancing for Adaptive Scientific Computations. , 2007, , .		85
143	Optimizing multiple queries on scientific datasets with partial replicas. , 2007, , .		4
144	Pathological Image Analysis Using the GPU: Stroma Classification for Neuroblastoma. , 2007, , .		12

#	ARTICLE	IF	CITATIONS
145	A Data Locality Aware Online Scheduling Approach for I/O-Intensive Jobs with File Sharing. Lecture Notes in Computer Science, 2007, , 141-160.	1.0	3
146	Dynamic Data-Driven Systems Approach for Simulation Based Optimizations. Lecture Notes in Computer Science, 2007, , 1213-1221.	1.0	1
147	Toward Optimizing Latency Under Throughput Constraints for Application Workflows on Clusters. Lecture Notes in Computer Science, 2007, , 173-183.	1.0	10
148	Scheduling File Transfers for Data-Intensive Jobs on Heterogeneous Clusters. Lecture Notes in Computer Science, 2007, , 214-223.	1.0	8
149	Development of Computer Based System To Aid Pathologists in Histological Grading of Follicular Lymphomas.. Blood, 2007, 110, 3318-3318.	0.6	2
150	MSSG: A Framework for Massive-Scale Semantic Graphs. , 2006, , .		2
151	Large Image Correction and Warping in a Cluster Environment. , 2006, , .		7
152	Hypergraph Partitioning for Automatic Memory Hierarchy Management. , 2006, , .		11
153	An XML-based System for Synthesis of Data from Disparate Databases. Journal of the American Medical Informatics Association: JAMIA, 2006, 13, 289-301.	2.2	8
154	Models, methods and middleware for grid-enabled multiphysics oil reservoir management. Engineering With Computers, 2006, 22, 349-370.	3.5	14
155	Data management and query--Hypergraph partitioning for automatic memory hierarchy management. , 2006, , .		6
156	Imaging and visual analysis--Large image correction and warping in a cluster environment. , 2006, , .		5
157	An extensible global address space framework with decoupled task and data abstractions. , 2006, , .		3
158	On Creating Efficient Object-relational Views of Scientific Datasets. , 2006, , .		0
159	Locality Conscious Processor Allocation and Scheduling for Mixed Parallel Applications. , 2006, , .		8
160	Servicing Seismic and Oil Reservoir Simulation Data Through Grid Data Services. Lecture Notes in Computer Science, 2006, , 129-142.	1.0	1
161	XML database support for distributed execution of data-intensive scientific workflows. SIGMOD Record, 2005, 34, 50-55.	0.7	9
162	Application of Grid-enabled technologies for solving optimization problems in data-driven reservoir studies. Future Generation Computer Systems, 2005, 21, 19-26.	4.9	25

#	ARTICLE	IF	CITATIONS
163	A simulation and data analysis system for large-scale, data-driven oil reservoir simulation studies. <i>Concurrency Computation Practice and Experience</i> , 2005, 17, 1441-1467.	1.4	14
164	Distributed Out-of-Core Preprocessing of Very Large Microscopy Images for Efficient Querying. , 2005, , .		5
165	Grid-based management of biomedical data using an XML-based distributed data management system. , 2005, , .		8
166	Tools for efficient subsetting and pipelined processing of large scale, distributed biomedical image data. <i>Advances in Parallel Computing</i> , 2005, , 403-422.	0.3	0
167	A Grid-Based Image Archival and Analysis System. <i>Journal of the American Medical Informatics Association: JAMIA</i> , 2005, 12, 286-295.	2.2	26
168	Towards Dynamic Data-Driven Optimization of Oil Well Placement. <i>Lecture Notes in Computer Science</i> , 2005, , 656-663.	1.0	10
169	A Scalable Parallel Graph Coloring Algorithm for Distributed Memory Computers. <i>Lecture Notes in Computer Science</i> , 2005, , 241-251.	1.0	23
170	Permuting Sparse Rectangular Matrices into Block-Diagonal Form. <i>SIAM Journal of Scientific Computing</i> , 2004, 25, 1860-1879.	1.3	96
171	The virtual microscope. <i>IEEE Transactions on Information Technology in Biomedicine</i> , 2003, 7, 230-248.	3.6	72
172	DATABASE SUPPORT FOR DATA-DRIVEN SCIENTIFIC APPLICATIONS IN THE GRID. <i>Parallel Processing Letters</i> , 2003, 13, 245-271.	0.4	40
173	Optimizing Reduction Computations In a Distributed Environment. , 2003, , .		6
174	Driving Scientific Applications by Data in Distributed Environments. <i>Lecture Notes in Computer Science</i> , 2003, , 355-364.	1.0	12
175	Processing large-scale multi-dimensional data in parallel and distributed environments. <i>Parallel Computing</i> , 2002, 28, 827-859.	1.3	47
176	Visualization of large data sets with the Active Data Repository. <i>IEEE Computer Graphics and Applications</i> , 2001, 21, 24-33.	1.0	43
177	Distributed processing of very large datasets with DataCutter. <i>Parallel Computing</i> , 2001, 27, 1457-1478.	1.3	161
178	A hypergraph-partitioning approach for coarse-grain decomposition. , 2001, , .		39
179	Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 1999, 10, 673-693.	4.0	375