

Thomas Neumann

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/376560/publications.pdf>

Version: 2024-02-01

115
papers

5,083
citations

430874

18
h-index

189892

50
g-index

128
all docs

128
docs citations

128
times ranked

1648
citing authors

#	ARTICLE	IF	CITATIONS
1	B2-Tree: Page-Based String Indexing in Concurrent Environments. Datenbank-Spektrum, 2022, 22, 11-22.	1.3	2
2	User-defined operators. Proceedings of the VLDB Endowment, 2022, 15, 1119-1131.	3.8	7
3	Efficient Evaluation of Arbitrarily-Framed Holistic SQL Aggregates and Window Functions. , 2022, , .		0
4	Adaptive Hybrid Indexes. , 2022, , .		3
5	Recursive SQL and GPU-support for in-database machine learning. Distributed and Parallel Databases, 2022, 40, 205-259.	1.6	1
6	Profiling dataflow systems on multiple abstraction levels. , 2021, , .		6
7	Building Advanced SQL Analytics From Low-Level Plan Operators. , 2021, , .		8
8	JSON Tiles: Fast Analytics on Semi-Structured Data. , 2021, , .		10
9	Self-Tuning Query Scheduling for Analytical Workloads. , 2021, , .		6
10	Tidy Tuples and Flying Start: fast compilation and fast execution of relational queries in Umbra. VLDB Journal, 2021, 30, 883-905.	4.1	20
11	To Partition, or Not to Partition, That is the Join Question in a Real System. , 2021, , .		17
12	TardisDB. , 2021, , .		4
13	In-Database Machine Learning with SQL on GPUs. , 2021, , .		3
14	ArrayQL for Linear Algebra within Umbra. , 2021, , .		1
15	A practical approach to groupjoin and nested aggregates. Proceedings of the VLDB Endowment, 2021, 14, 2383-2396.	3.8	3
16	Make the most out of your SIMD investments: counter control flow divergence in compiled query pipelines. VLDB Journal, 2020, 29, 757-774.	4.1	6
17	Building blocks for persistent memory. VLDB Journal, 2020, 29, 1223-1241.	4.1	17
18	Rethinking Logging, Checkpoints, and Recovery for High-Performance Storage Engines. , 2020, , .		19

#	ARTICLE	IF	CITATIONS
19	Concurrent online sampling for all, for free. , 2020, , .		6
20	RadixSpline. , 2020, , .		61
21	Benchmarking learned indexes. Proceedings of the VLDB Endowment, 2020, 14, 1-13.	3.8	56
22	Tree-Encoded Bitmaps. , 2020, , .		6
23	Scalable and robust latches for database systems. , 2020, , .		8
24	ARTful Skyline Computation for In-Memory Database Systems. Communications in Computer and Information Science, 2020, , 3-12.	0.5	0
25	Freedom for the SQL-Lambda. , 2020, , .		4
26	MLearn. , 2019, , .		5
27	Persistent Memory I/O Primitives. , 2019, , .		46
28	Versioning in Main-Memory Database Systems. , 2019, , .		4
29	On the Impact of Memory Allocation on High-Performance Query Processing. , 2019, , .		18
30	Scalable Analytics on Fast Data. ACM Transactions on Database Systems, 2019, 44, 1-35.	2.8	6
31	Estimating Cardinalities with Deep Sketches. , 2019, , .		22
32	Scalable garbage collection for in-memory MVCC systems. Proceedings of the VLDB Endowment, 2019, 13, 128-141.	3.8	24
33	Query optimization through the looking glass, and what we found running the Join Order Benchmark. VLDB Journal, 2018, 27, 643-668.	4.1	74
34	Everything you always wanted to know about compiled and vectorized queries but were afraid to ask. Proceedings of the VLDB Endowment, 2018, 11, 2209-2222.	3.8	31
35	Adaptive Execution of Compiled Queries. , 2018, , .		31
36	Make the most out of your SIMD investments. , 2018, , .		5

#	ARTICLE	IF	CITATIONS
37	Adaptive Optimization of Very Large Join Queries. , 2018, , .		30
38	Managing Non-Volatile Memory in Database Systems. , 2018, , .		75
39	Get Real. , 2018, , .		30
40	LeanStore: In-Memory Data Management beyond Main Memory. , 2018, , .		48
41	Order Indexes: supporting highly dynamic hierarchical data in relational main-memory database systems. VLDB Journal, 2017, 26, 55-80.	4.1	2
42	Automatic algorithm transformation for efficient multi-snapshot analytics on temporal graphs. Proceedings of the VLDB Endowment, 2017, 10, 877-888.	3.8	17
43	Efficient Batched Distance, Closeness and Betweenness Centrality Computation in Unweighted and Weighted Graphs. Datenbank-Spektrum, 2017, 17, 169-182.	1.3	4
44	HyPerInsight. , 2017, , .		6
45	Evaluation of parallel graph loading techniques. , 2016, , .		1
46	Data Blocks. , 2016, , .		93
47	High-Performance Geospatial Analytics in HyPerSpace. , 2016, , .		18
48	Flow-Join: Adaptive skew handling for distributed joins over high-speed networks. , 2016, , .		37
49	The ART of practical synchronization. , 2016, , .		57
50	Scaling HTM-Supported Database Transactions to Many Cores. IEEE Transactions on Knowledge and Data Engineering, 2016, 28, 297-310.	5.7	11
51	High-speed query processing over high-speed networks. Proceedings of the VLDB Endowment, 2015, 9, 228-239.	3.8	43
52	How good are query optimizers, really?. Proceedings of the VLDB Endowment, 2015, 9, 204-215.	3.8	305
53	Indexing highly dynamic hierarchical data. Proceedings of the VLDB Endowment, 2015, 8, 986-997.	3.8	8
54	Fast Serializable Multi-Version Concurrency Control for Main-Memory Database Systems. , 2015, , .		138

#	ARTICLE	IF	CITATIONS
55	Benchmarking Elastic Query Processing on Big Data. Lecture Notes in Computer Science, 2015, , 37-44.	1.3	0
56	High-Performance Main-Memory Database Systems and Modern Virtualization. , 2015, , .		3
57	Supporting hierarchical data in SAP HANA. , 2015, , .		9
58	Massively Parallel NUMA-Aware Hash Joins. Lecture Notes in Computer Science, 2015, , 3-14.	1.3	11
59	Scaling Up Mixed Workloads: A Battle of Data Freshness, Flexibility, and Scheduling. Lecture Notes in Computer Science, 2015, , 97-112.	1.3	5
60	Engineering high-performance database engines. Proceedings of the VLDB Endowment, 2014, 7, 1734-1741.	3.8	5
61	Morsel-driven parallelism. , 2014, , .		166
62	The more the merrier. Proceedings of the VLDB Endowment, 2014, 8, 449-460.	3.8	58
63	Heterogeneity-conscious parallel query execution. , 2014, , .		17
64	On-the-fly token similarity joins in relational databases. , 2014, , .		4
65	The linked data benchmark council. SIGMOD Record, 2014, 43, 27-31.	1.2	56
66	One DBMS for all. , 2014, , .		11
67	HyPer Beyond Software: Exploiting Modern Hardware for Main-Memory Database Systems. Datenbank-Spektrum, 2014, 14, 173-181.	1.3	2
68	Main-memory database systems. , 2014, , .		3
69	Exploiting hardware transactional memory in main-memory databases. , 2014, , .		74
70	Locality-sensitive operators for parallel main-memory database clusters. , 2014, , .		31
71	TPC-H Analyzed: Hidden Messages and Lessons Learned from an Influential Benchmark. Lecture Notes in Computer Science, 2014, , 61-76.	1.3	74
72	The Database Group at TUM. SIGMOD Record, 2014, 43, 55-60.	1.2	0

#	ARTICLE	IF	CITATIONS
73	The Linked Data Benchmark Council Project. Datenbank-Spektrum, 2013, 13, 121-129.	1.3	13
74	CPU and cache efficient management of memory-resident databases. , 2013, , .		22
75	ScyPer. , 2013, , .		18
76	DeltaNI. , 2013, , .		9
77	Instant loading for main memory databases. Proceedings of the VLDB Endowment, 2013, 6, 1702-1713.	3.8	51
78	The adaptive radix tree: ARTful indexing for main-memory databases. , 2013, , .		242
79	Elasticity in Cloud Databases and Their Query Processing. International Journal of Data Warehousing and Mining, 2013, 9, 1-20.	0.6	3
80	Compacting transactional data in hybrid OLTP&OLAP databases. Proceedings of the VLDB Endowment, 2012, 5, 1424-1435.	3.8	32
81	Fast approximation of steiner trees in large graphs. , 2012, , .		8
82	The mainframe strikes back. , 2012, , .		3
83	Massively parallel sort-merge joins in main memory multi-core database systems. Proceedings of the VLDB Endowment, 2012, 5, 1064-1075.	3.8	132
84	Effective and Robust Pruning for Top-Down Join Enumeration Algorithms. , 2012, , .		11
85	HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. , 2011, , .		384
86	Efficiently compiling efficient query plans for modern hardware. Proceedings of the VLDB Endowment, 2011, 4, 539-550.	3.8	309
87	The mixed workload CH-benCHmark. , 2011, , .		82
88	How to efficiently snapshot transactional data. , 2011, , .		12
89	Characteristic sets: Accurate cardinality estimation for RDF queries with multiple joins. , 2011, , .		132
90	One Size Fits all, Again! The Architecture of the Hybrid OLTP&OLAP Database Management System HyPer. Lecture Notes in Business Information Processing, 2011, , 7-23.	1.0	5

#	ARTICLE	IF	CITATIONS
91	Accelerating queries with group-by and join by groupjoin. Proceedings of the VLDB Endowment, 2011, 4, 843-851.	3.8	19
92	x-RDF-3X. Proceedings of the VLDB Endowment, 2010, 3, 256-263.	3.8	96
93	The RDF-3X engine for scalable management of RDF data. VLDB Journal, 2010, 19, 91-113.	4.1	478
94	Active knowledge. , 2010, , .		26
95	Scalable join processing on very large RDF graphs. , 2009, , .		146
96	Preventing bad plans by bounding the impact of cardinality estimation errors. Proceedings of the VLDB Endowment, 2009, 2, 982-993.	3.8	92
97	Generating optimal DAG-structured query evaluation plans. Computer Science - Research and Development, 2009, 24, 103-117.	2.7	2
98	Distributed top-k aggregation queries at large. Distributed and Parallel Databases, 2009, 26, 3-27.	1.6	14
99	Query simplification. , 2009, , .		21
100	ANGIE. Proceedings of the VLDB Endowment, 2009, 2, 1570-1573.	3.8	7
101	System R (R*) Optimizer. , 2009, , 2900-2905.		0
102	Query Optimization (in Relational Databases). , 2009, , 2273-2278.		1
103	Exploiting social relations for query expansion and result ranking. , 2008, , .		52
104	Faster Join Enumeration for Complex Queries. , 2008, , .		1
105	Dynamic programming strikes back. , 2008, , .		61
106	Near-optimal dynamic replication in unstructured peer-to-peer networks. , 2008, , .		19
107	Efficient top-k querying over social-tagging networks. , 2008, , .		106
108	Social recommendations at work. , 2008, , .		0

#	ARTICLE	IF	CITATIONS
109	RDF-3X. Proceedings of the VLDB Endowment, 2008, 1, 647-659.	3.8	370
110	Smooth Interpolating Histograms with Error Guarantees. Lecture Notes in Computer Science, 2008, , 126-138.	1.3	10
111	Optimizing Distributed Top-k Queries. Lecture Notes in Computer Science, 2008, , 337-349.	1.3	1
112	Making SENSE. Proceedings of the VLDB Endowment, 2008, 1, 1480-1483.	3.8	5
113	A time machine for text search. , 2007, , .		64
114	Algebraic query optimization for distributed top-k queries. Computer Science - Research and Development, 2007, 21, 197-211.	0.9	2
115	Indexing Set-Valued Attributes with a Multi-level Extendible Hashing Scheme. Lecture Notes in Computer Science, 2007, , 98-108.	1.3	3