

STRATOS IDREOS

Professor in Computer Science

Ph.D.: June 2010

Harvard John A. Paulson School of Engineering and Applied Sciences

Citations > 6600

stratos.seas.harvard.edu stratos@seas.harvard.edu daslab.seas.harvard.edu

Hindex:42, i10:86

**MAPPING DATA SYSTEMS DNA - SELF-DESIGNING DATA SYSTEMS - HARDWARE/SOFTWARE CO-DESIGN
ACCELERATING DATA SCIENCE - DATA EXPLORATION - END-TO-END DATA SCIENCE PERFORMANCE**

AWARDS



2023 Sloan Research Fellowship



2023 Capers W. McDonald and Marion K. McDonald Award

For Excellence in Mentoring and Advising.



2022 ACM SIGMOD Test of time Award

For the ACM SIGMOD'12 paper on "NoDB: efficient query execution on raw data files".



2020 ACM SIGMOD Contributions Award

For "the innovative work in the data management community to encourage scientific reproducibility of our publications".



2019 Department of Energy Early Career Award. For research on Data Structure Alchemy



2017 Best of ACM SIGMOD

Our paper on "Monkey: Optimal Navigable Key-value Store" was selected as one of the (4) best papers of the ACM SIGMOD International Conference on Management of Data



2017 & 2018 ACM SIGMOD Most Reproducible Paper Award



2016 Facebook Faculty Award. For our work on self-designing key-value stores



2016 & 2018 NetApp Faculty Award. For our work on self-designing data systems



2015 Rising Star Award. By the IEEE Technical Committee on Data Engineering.

"For breakthroughs in adaptive data management systems that allow users rich data exploration capabilities without a priori knowledge of the data or queries".



2015 NSF Career Award. For research on Evolutionary Data Systems.


2015 Communications of the ACM Research Highlights

For the ACM SIGMOD'12 paper on NoDB


2012 Veni Award. By the Dutch Organization For Research (NWO).

Awarded to "the most promising young researchers in the Netherlands".



2012 Best of VLDB. Our paper on Concurrency Control for Adaptive Indexing was selected as one of the best papers in the 38th International Conference on Very Large Databases (VLDB).

2011 Cor Baayen Award. European Research Consortium for Informatics and Mathematics (ERCIM).

Named "Most promising European young researcher in computer science and applied mathematics".



Best Ph.D. thesis across all areas of computer science and applied mathematics over all 'ERCIM countries': Austria, Belgium, Cyprus, Czech Republic, Finland, France, Germany, Greece, Hungary, Italy, Luxembourg, Norway, Poland, Portugal, Spain, Sweden, Switzerland, The Netherlands, UK.


2011 ACM SIGMOD Jim Gray Doctoral Dissertation Award

By the ACM Special Interest Group on Management of Data


2011 VLDB Best Paper Award in Challenges and Visions

2010 IBM zEnterprise System Recognition Award

For contributions in the IBM ISAO Smart Analytics Big Data product.


2010 Named "Distinguished Scientist Excelling in Research Abroad"

By the Hellenic Ministry of National Defense


2005 Best Student Paper Award

In 9th European Conf. on Research and Advanced Technology for Digital Libraries (ECDL)

Student Awards

1st Place ACM SIGMOD 2022 Undergraduate Research Competition. Winner: Joshua Pan

3rd Place ACM SIGMOD 2021 Undergraduate Research Competition. Winner: Demi Guo

1st Place ACM SIGMOD 2020 Undergraduate Research Competition. Winner: Meena Jagadeesan

2nd Place ACM SIGMOD 2020 Undergraduate Research Competition. Winner: Neil Band

1st Place ACM SIGMOD 2019 Undergraduate Research Competition. Winner: H. Gupta, V. Jain, J. Lennon

1st Place ACM SIGMOD 2018 Undergraduate Research Competition. Winner: T. Lively, C. Mendizabal, L. Schroeder

1st Place ACM SIGMOD 2017 Undergraduate Research Competition. Winner: Graham Lustiber

1st Place ACM SIGMOD 2016 Undergraduate Research Competition. Winner: Alex Liu

3rd Place ACM SIGMOD 2016 Undergraduate Research Competition. Winner: Wilson Qin

3rd Place ACM SIGMOD 2015 Undergraduate Research Competition. Winner: Stella Pantela

Microsoft PhD Fellowship 2019. Winner: Wilson Qin

Microsoft PhD Fellowship 2016. Winner: Lukas Maas

EMPLOYMENT

Gordon McKay Professor of Computer Science, Harvard University , USA	1/2023 - now
Founder, CTO, Leibniz Data Systems , USA	1/2021 - now
Associate Professor in Computer Science, Harvard University , USA	1/2019 - 12/2022
Principal AI Software Architect SimpleBet , USA	11/2018 - 4/2020
Assistant Professor in Computer Science, Harvard University , USA	1/2014 - 12/2018
Scientific Staff Member (Tenure-track) Dutch National Research Center for Mathematics and Computer Science, CWI , The Netherlands	4/2010 - 12/2013
Visiting Professor, Data Intensive Applications and Systems group, EPFL , Switzerland	9/2013 - 12/2013
Visiting Professor, Database group, University of Trento , Trento, Italy	11/2011
Visiting Professor, Database group, National University of Singapore , Singapore	4/2011
Research Intern, Database group, IBM Almaden Research Center , USA	7/2010 - 12/2010
Scientific Staff, Data Intensive Applications and Systems group, EPFL , Switzerland	12/2009 - 4/2010
Research Intern, Database group, Microsoft Research Redmond, USA	1/2009 - 4/2009
Researcher (Ph.D.), Database Architectures group, Dutch National Research Center for Mathematics and Computer Science, CWI , The Netherlands	10/2005 - 11/2009
Researcher, Intelligent Systems Laboratory, Technical University of Crete , Greece	8/2003 - 9/2005

EDUCATION

Ph.D. in Computer Science , University of Amsterdam, The Netherlands Thesis: Database Cracking: Towards Auto-Tuning Database Kernels Supervised by Prof. Martin Kersten and Prof. Stefan Manegold	10/2005 - 6/2010
Master in Computer Engineering , Technical University of Crete, Greece Thesis: Distributed Evaluation of Continuous Equi-join Queries over Large Structured Overlay Networks Supervised by Prof. Manolis Koubarakis	7/2003 - 9/2005
Diploma in Electronic and Computer Engineering Technical University of Crete, Greece (5 year school) Thesis: P2P-DIET: A query and notification service based on mobile agents for rapid implementation of peer-to-peer applications Supervised by Prof. Manolis Koubarakis	9/1997 - 6/2003

STUDENT SUPERVISION & TEACHING

Brian Hentschel, Sanket Purandere, Lukas Maas, Abdul Wasay, Michael Kester, Wilson Qin	PhD
Konstantinos Zoumpatianos, Niv Dayan, Manos Athanassoulis, Subarna Chatterjee, Siqiang Luo, Utku Sirin, Hao Jiang	Postdoc
Demi Guo, Meena Jagadeesan, Kyle Deeds, Neil Band, Graham Lustiber, Drhuv Gupta, Tarik Moon, Stella Pantella, Alex Liu, Ore Babarinsa, Yifan Wu	Undergrad

Data Systems	Spring 2014, Fall 2015-2024	http://daslab.seas.harvard.edu/classes/cs165/
Big Data Systems	Fall 2014, Spring 2016-2023	http://daslab.seas.harvard.edu/classes/cs265/

PUBLICATIONS

REFEREED CONFERENCES and JOURNALS

- 1 Limousine: Blending Learned and Classical Indexes to Self-Design Larger-than-Memory Cloud Storage Engines. Subarna Chatterjee, Mark F. Pekala, Lev Kruglyak, and Stratos Idreos. In Proceedings of the ACM Management of Data 2, 1, Article 47 (February 2024), (SIGMOD), 2024 **SIGMOD 2024**
- 2 The Image Calculator: 10x Faster Image-AI Inference by Replacing JPEG with Self-designing Storage Format. Utku Sirin and Stratos Idreos. In Proceedings of the ACM Management of Data 2, 1, Article 52 (February 2024), (SIGMOD), 2024 **SIGMOD 2024**
- 3 Data Structures for Data-Intensive Applications: Tradeoffs and Design Guidelines. Manos Athanassoulis, Stratos Idreos, Dennis E. Shasha. Foundations and Trends in Databases 13(1-2): 1-168 (2023) **FNT 2023**
- 4 Adaptive Indexing of Objects with Spatial Extent. Fatemeh Zardbani, Nikos Mamoulis, Stratos Idreos, and Panagiotis Karras. In Proceedings of the Very Large Databases Endowment (PVLDB), 16(9): 2248-2260 (2023) **PVLDB 2023**
- 5 Real-Time LSM-Trees for HTAP Workloads. Hemant Saxena, Lukasz Golab, Stratos Idreos, and Ihab F. Ilyas. In Proceedings of the IEEE International Conference on Data Engineering (ICDE), 2023 **ICDE 2023**
- 6 The Case for Distributed Shared-Memory Databases with RDMA-Enabled Memory Disaggregation. Ruihong Wang, Jianguo Wang, Stratos Idreos, M. Tamer Özsu, Walid G. Aref. In Proceedings of the Very Large Databases Endowment (PVLDB), 16(1): 15-22 (2022) **PVLDB 2022**
- 7 Cosine: A Cloud-Cost Optimized Self-Designing Key-Value Storage Engine. Subarna Chatterjee, Meena Jagadeesan, Wilson Qin, and Stratos Idreos. In Proceedings of the Very Large Databases Endowment (PVLDB), 2022 **PVLDB 2022**
- 8 SNARF: A Learning-Enhanced Range Filter. Kapil Vaidya, Subarna Chatterjee, Eric Knorr, Michael Mitzenmacher, Stratos Idreos, Tim Kraska. In Proceedings of the Very Large Databases Endowment (PVLDB), 2022 **PVLDB 2022**
- 9 Entropy Learned Hashing: 10X Faster Hashing with Controllable Uniformity. Brian Hentschel, Utku Sirin, and Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2022 **SIGMOD 2022**
- 10 Proteus: A Self-Designing Range Filter. Eric Knorr, Baptiste Lemaire, Andrew Lim, Huanchen Zhang, Siqiang Luo, Stratos Idreos, Michael Mitzenmacher. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2022 **SIGMOD 2022**
- 11 More or Less: When and How to Build Convolutional Neural Network Ensembles. Abdul Wasay and Stratos Idreos. In Proceedings of the International Conference on Learning Representations (ICLR), 2021 **ICLR 2021**
- 12 Stacked Filters: Learning to Filter by Structure. Kyle Deeds, Brian Hentschel, and Stratos Idreos. In Proceedings of the VLDB Endowment, vol. 14, no. 4, pp. 600 - 612, 2021. **PVLDB 2021**
- 13 Rosetta: A Robust Space-Time Optimized Range Filter for Key-Value Stores. Siqiang Luo, Subarna Chatterjee, Rafael Ketsetsidis, Niv Dayan, Wilson Qin, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2020 **SIGMOD 2020**
- 14 MotherNets: Rapid Deep Ensemble Learning. Abdul Wasay, Brian Hentschel, Yuze Liao, Sanyuan Chen, Stratos Idreos. The Conference on Machine Learning and Systems, 2020 **MLSys 2020**
- 15 Optimal Column Layout for Hybrid Workloads. Manos Athanassoulis, Kenneth Bogh, Stratos Idreos. In Proceedings of the Very Large Databases Endowment (PVLDB), 2019 **PVLDB 2019**
- 16 Learning Data Structure Alchemy. Stratos Idreos, Kostas Zoumpatianos, Subarna Chatterjee, Wilson Qin, Abdul Wasay, Brian Hentschel, Mike Kester, Niv Dayan, Demi Guo, Minseo Kang, Yiyu Sun. IEEE Data Engineering Bull. Sep, 2019 **IEEE DEBull 2019**
- 17 The Log-Structured Merge-Bush & the Wacky Continuum. Niv Dayan, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2019 **SIGMOD 2019**

Stratos Idreos, Curriculum Vitae, January 2024

18	Design Continuums and the Path Toward Self-Designing Key-Value Stores that Know and Learn. Stratos Idreos, Niv Dayan, Wilson Qin, Mali Akmanalp, Sophie Hilgard, Andrew Ross, James Lennon, Varun Jain, Harshita Gupta, David Li, and Zichen Zhu. Proceedings of CIDR Conference on Innovative Data Systems Research, 2019.	CIDR 2019
19	Optimal Bloom Filters and Adaptive Merging for LSM-Trees. Niv Dayan, Manos Athanassoulis, Stratos Idreos. ACM Transactions on Database Systems, Dec 2018	TODS 2018
20	The Periodic Table of Data Structures. Stratos Idreos, Kostas Zoumpatianos, Manos Athanassoulis, Niv Dayan, Brian Hentschschel, Michael S. Kester, Demi Guo, Lukas Maas, Wilson Qin, Abdul Wasay, Yiyu Sun. IEEE Data Engineering Bull. Sep, 2018	IEEE DEBull 2018
21	Smooth Scan: Robust Access Path Selection Without Cardinality Estimation. Renata Borovica, Stratos Idreos, Anastassia Ailamaki, Marcin Zukowski, Campbell Fraser. The International Journal on Very Large Databases (VLDBJ), 2018	VLDBJ 2018
22	The Data Calculator: Data Structure Design and Cost Synthesis From First Principles, and Learned Cost Models. Stratos Idreos, Konstantinos Zoumpatianos, Brian Hentschel, Michael Kester, Demi Guo. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2018	SIGMOD 2018
23	Dostoevsky: Better Space-Time Trade-Offs for LSM-Tree Based Key-Value Stores via Adaptive Removal of Superfluous Merging. Niv Dayan, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2018	SIGMOD 2018
24	Column Sketches: A Scan Accelerator for Rapid and Robust Predicate Evaluation. Brian Hentschel, Michael Kester, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2018	SIGMOD 2018
25	Monkey: Optimal Navigable Key-Value Store. Niv Dayan, Manos Athanassoulis, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2017 Best of SIGMOD.	SIGMOD 2017
26	Access Path Selection in Main-Memory Optimized Data Systems: Should I Scan or Should I Probe? Michael Kester, Manos Athanassoulis, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2017	SIGMOD 2017
27	Data Canopy: Accelerating Exploratory Statistical Analysis. Abdul Wasay, Xinding Wei, Niv Dayan, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2017	SIGMOD 2017
28	UpBit: Scalable In-Memory Updatable Bitmap Indexing. Manos Athanassoulis, Zheng Yan, Stratos Idreos. In Proc. of the ACM SIGMOD Intern. Conference on Management of Data, 2016.	SIGMOD 2016
29	GeckoFTL: Scalable Flash Translation Techniques For Very Large Flash Devices. Niv Dayan, Phillippe Bonnet, Stratos Idreos. In Proc. of the ACM SIGMOD Intern. Conference on Management of Data, 2016	SIGMOD 2016
30	Adaptive Indexing over Encrypted Data. Panagiotis Karras, Saad Malhotra, Rudrika Bhatt, Artyom Nikitin, Denis Antyukhov, and Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2016	SIGMOD 2016
31	Designing Access Methods: The RUM Conjecture. Manos Athanassoulis, Michael Kester, Lukas Maas, Radu Stoica, Stratos Idreos, Anastassia Ailamaki, and Mark Callaghan. In Proceedings of the International Conference on Extending Database Technology (EDBT), 2016	EDBT 2016
32	ADS: The Adaptive Data Series Index. Konstantinos Zoumpatianos, Stratos Idreos, Themis Palpanas. The International Journal on Very Large Databases (VLDBJ), 2016	VLDBJ 2016
33	Queriosity: Automated Data Exploration. Abdul Wasay, Manos Athanassoulis, and Stratos Idreos. In Proceedings of the IEEE International Congress on Big Data, 2015	BIG DATA 2016
34	Holistic Indexing in Main-memory Column-stores. Eleni Petraki, Stratos Idreos, Stefan Manegold. In Proc. of the ACM SIGMOD Intern. Conference on Management of Data, 2015	SIGMOD 2015
35	Smooth Scan: Statistics-Oblivious Access Paths. Renata Borovica, Stratos Idreos, Anastassia Ailamaki, Marcin Zukowski, Campbell Fraser. In Proceedings of the IEEE International Conference on Data Engineering (ICDE), 2015	ICDE 2015

Stratos Idreos, Curriculum Vitae, January 2024

36	NoDB: Efficient Query Execution on Raw Data Files. Ioannis Alagiannis, Renata Borovica, Miguel Branco, Stratos Idreos and Anastasia Ailamaki. Communications of the ACM, Research Highlights, 2015	CACM 2015
37	H2O: A Hands-free Adaptive Store. Ioannis Alagiannis, Stratos Idreos, Anastasia Ailamaki. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2014	SIGMOD 2014
38	Indexing for Interactive Exploration of Big Data Series. Konstantinos Zoumpatianos, Stratos Idreos, Themis Palpanas. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2014	SIGMOD 2014
39	Joins On Encoded and Partitioned Data. Jae-Gil Lee, Gopi Attaluri, Ronald Barber, Naresh Chainani, Oliver Draese, Frederick Ho, Stratos Idreos, Min-Soo Kim, Sam Lightstone, Guy Lohman, Konstantinos Morfonios, Keshava Murthy, Ippokratis Pandis, Lin Qiao, Vijayshankar Raman, Vincent Kulandai Samy, Richard Sidle, Knut Stolze and Liping Zhang. In Proceedings of the Very Large Databases Endowment (PVLDB), 2014	PVLDB 2014
40	Transactional Support for Adaptive Indexing. Goetz Graefe, Felix Halim, Stratos Idreos, Harumi Kuno, Stefan Manegold, Bernhard Seeger. The International Journal on Very Large Databases (VLDBJ). 23(2):303-328, 2014	VLDBJ 2014
41	Transactions on Large-Scale Data- and Knowledge-Centered Systems. Christos Tryfonopoulos, Stratos Idreos, Manolis Koubarakis, Paraskevi Raftopoulou. Distributed Large-Scale Information Filtering. XIII Lecture Notes in Computer Science. 13:91-122, 2014	TLSDKS 2014
42	dbTouch: Analytics at your Fingertips. Stratos Idreos and Erietta Liarou. In Proceedings of the 7th International Conference on Innovative Data Systems Research (CIDR), 2013	CIDR 2013
43	Enhanced Stream Processing in a DBMS Kernel. Erietta Liarou, Stratos Idreos, Stefan Manegold and Martin L. Kersten. In Proceedings of the 16th International Conference on Extending Database Technology (EDBT), 2013	EDBT 2013
44	The Design and Implementation of Modern Column-Oriented Database Systems. Daniel Abadi, Peter Boncz, Stavros Harizopoulos, Stratos Idreos, Samuel Madden. Foundations and Trends in Databases. 5(3):197-280, 2013	FNT 2013
45	NoDB: Efficient Query Execution on Raw Data Files. Ioannis Alagiannis, Renata Borovica, Miguel Branco, Stratos Idreos and Anastasia Ailamaki. In Proceedings of the 32nd ACM SIGMOD International Conference on Management of Data, 2012	SIGMOD 2012
46	Stochastic Database Cracking: Towards Robust Adaptive Indexing in Main-Memory Column-Stores. Felix Halim, Stratos Idreos, Panagiotis Karras and Roland H. C. Yap. In Proceedings of the Very Large Databases Endowment (PVLDB), 5(6): 502-513, 2012	PVLDB 2012
47	Concurrency Control for Adaptive Indexing. Goetz Graefe, Felix Halim, Stratos Idreos, Harumi A. Kuno, Stefan Manegold. In Proceedings of the Very Large Databases Endowment (PVLDB), 5(7): 656-667, 2012. Best of VLDB .	PVLDB 2012
48	Business Analytics in (a) Blink. Ronald Barber, Peter Bendel, Marco Czech, Oliver Draese, Frederick Ho, Namik Hrlé, Stratos Idreos, Min-Soo Kim, Oliver Koeth, Jae-Gil Lee, Tianchao Tim Li, Guy M. Lohman, Konstantinos Morfonios, René Müller, Keshava Murthy, Ippokratis Pandis, Lin Qiao, Vijayshankar Raman, Richard Sidle, Knut Stolze, Sandor Szabo. IEEE Data Engineering Bull. 35(1): 9-14, 2012	IEEE DEBull 2012
49	MonetDB: Two Decades of Research in Column-oriented Database Architectures. Stratos Idreos, Fabian Groffen, Niels Nes, Stefan Manegold, K. Sjoerd Mullender, Martin L. Kersten. IEEE Data Engineering Bull. 35(1): 40-45, 2012	IEEE DEBull 2012
50	Too Many Links in the Horizon; What is next? Linked Views and Linked History. Erietta Liarou and Stratos Idreos. In Proc. of the 10th International Semantic Web Conference (ISWC), 2011	ISWC 2011
51	Here are my Data Files. Here are my Queries. Where are my Results? Stratos Idreos, Ioannis Alagiannis, Ryan Johnson and Anastasia Ailamaki. In Proceedings of the 5th International Conference on Innovative Data Systems Research (CIDR), 2011	CIDR 2011
52	Merging What's Cracked, Cracking What's Merged: Adaptive Indexing in Main-Memory Column-Stores. Stratos Idreos, Stefan Manegold, Harumi Kuno and Goetz Graefe. In Proceedings of the Very Large Databases Endowment (PVLDB), 4(9): 585-597, 2011	PVLDB 2011

Stratos Idreos, Curriculum Vitae, January 2024

- 53 The Researcher's Guide to the Data Deluge: Querying a Scientific Database in Just a Few Seconds. Martin Kersten, Stratos Idreos, Stefan Manegold and Erietta Liarou. In Proceedings of the Very Large Databases Endowment (PVLDB), 4(12): 174-177, 2011 **PVLDB 2011**
Best paper award in Challenges and Visions
- 54 Benchmarking Adaptive Indexing. Goetz Graefe, Stratos Idreos, Harumi Kuno and Stefan Manegold. In Proceedings of the 2nd TPC Technology Conference on Performance Evaluation & Benchmarking (TPCTC), 2010 **TPCTC 2010**
- 55 Estimating the Compression Fraction of an Index Using Sampling. Stratos Idreos, Raghav Kaushik, Vivek Narasayya, Ravishankar Ramamurthy. In Proc. of the 22nd IEEE Intern. Conference in Data Engineering (ICDE), 2010 **ICDE 2010**
- 56 Self-organizing Tuple Reconstruction In Column-stores. Stratos Idreos, Martin Kersten and Stefan Manegold. In Proc. of the 29th ACM SIGMOD International Conference on Management of Data, 2009 **SIGMOD 2009**
- 57 Exploiting the Power of Relational Databases for Efficient Stream Processing. Erietta Liarou, Romulo Goncalves and Stratos Idreos. In Proc. of the 19th International. Conference on Extending Database Technology (EDBT), 2009 **EDBT 2009**
- 58 Continuous Multi-way Joins over Distributed Hash Tables. Stratos Idreos, Erietta Liarou and Manolis Koubarakis. In Proc. of the 11th Intern. Conference on Extending Database Technology (EDBT), 2008 **EDBT 2008**
- 59 Continuous RDF Query Processing over DHTs. Erietta Liarou, Stratos Idreos and Manolis Koubarakis. In Proceedings of the 6th International Semantic Web Conference (ISWC), 2007 **ISWC 2007**
- 60 Updating a Cracked Database. Stratos Idreos, Martin Kersten and Stefan Manegold. In Proceedings of the 27th ACM SIGMOD International Conference on Management of Data, 2007 **SIGMOD 2007**
- 61 Database Cracking. Stratos Idreos, Martin Kersten and Stefan Manegold. In Proceedings of the 3rd International Conference on Innovative Data Systems Research (CIDR), 2007 **CIDR 2007**
- 62 Evaluating Conjunctive Triple Pattern Queries over Large Structured Overlay Networks. Erietta Liarou, Stratos Idreos, Manolis Koubarakis. In Proc. of the 5th Intern. Semantic Web Conference (ISWC), 2006 **ISWC 2006**
- 63 Distributed Evaluation of Continuous Equi-join Queries over Large Structured Overlay Networks. Stratos Idreos, Christos Tryfonopoulos and Manolis Koubarakis. In Proc. of the 22nd IEEE International Conference in Data Engineering (ICDE), 2006 **ICDE 2006**
- 64 Publish/Subscribe Functionality in IR Environments using Structured Overlay Networks. Christos Tryfonopoulos, Stratos Idreos and Manolis Koubarakis. In Proceedings of the 28th Annual International ACM SIGIR Conference, 2005 **SIGIR 2005**
- 65 LibraRing: An Architecture for Future Digital Libraries Based on Structured Overlay Networks. Christos Tryfonopoulos, Stratos Idreos and Manolis Koubarakis. In Proceedings of the 9th European Conference on Research and Advanced Technology for Digital Libraries (ECDL), 2005 **ECDL 2005**
Best student paper award
- 66 Publish/Subscribe for RDF-based P2P Networks. Paul-Alexandru Chirita, Stratos Idreos, Manolis Koubarakis and Wolfgang Nejdl. In Proc. of the 1st European Semantic Web Conference (ESWC), 2004 **ESWC 2004**
- 67 Selective Information Dissemination in P2P Networks: Problems and Solutions. Manolis Koubarakis, Christos Tryfonopoulos, Stratos Idreos and Yannis Drougas. ACM SIGMOD Record, Special issue on Peer-to- Peer Data Management, 32(3), September 2003 **SIGMOD Record 2003**

TUTORIALS

- 1 Deep Learning: Systems and Responsibility. Abdul Wasay, Subarna Chatterjee, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2021 **SIGMOD 2021**
- 2 Key-Value Storage Engines. Stratos Idreos, Mark Callaghan. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2020 **SIGMOD 2020**
- 3 From Auto-tuning One Size Fits All to Self-designed and Learned Data-intensive Systems. Stratos Idreos, Tim Kraska. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2019 **SIGMOD 2019**

Stratos Idreos, Curriculum Vitae, January 2024

4	Design Tradeoffs of Data Access Methods. Manos Athanassoulis, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2016	SIGMOD 2016
5	Overview of Data Exploration Techniques. Stratos Idreos, Olga Papaemmanouil, Surajit Chaudhuri. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2015	SIGMOD 2015
6	Adaptive Indexing in Modern Database Kernels. Stratos Idreos, Stefan Manegold and Goetz Graefe. In Proceedings of the 9th International Conference on Extending Database Technology (EDBT), 2012	EDBT 2012

SELECTED REFEREED SHORT PAPERS

1	The Automatic Scientist will be a Data System. Stratos Idreos. In Proceedings of the Biennial Conference on Innovative Data Systems Research (CIDR), 2017	CIDR 2017
2	Main Memory Adaptive Denormalization. Alex Liu, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2016	SIGMOD 2016
3	Adaptive Data Skipping in Main-Memory Systems. Wilson Qin, Stratos Idreos. In Proceedings of the ACM SIGMOD Intern. Conference on Management of Data, 2016	SIGMOD 2016
4	RINSE: Interactive Data Series Exploration with ADS+. Konstantinos Zoumpatianos, Stratos Idreos, Themis Palpanas. In Proceedings of the Very Large Databases Endowment (PVLDB), 2015	PVLDB 2016
5	One Loop Does Not Fit All. Styliani Pantella and Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2015	SIGMOD 2015
6	JAFAR: Near-Data Processing for Databases. Ore Babarinsa and Stratos Idreos. In Proceedings of the ACM SIGMOD Intern. Conference on Management of Data, 2015	SIGMOD 2015
7	Beyond the Wall: Near-Data Processing for Databases. Sam Xi, Ore Babarinsa, Manos Athanassoulis, and Stratos Idreos. In Proceedings of the 11th International Workshop on Data Management on New Hardware (DaMoN), 2015	DaMoN 2015
8	Database Cracking: Fancy Scan, not Poor Man Sort! Holger Pirk, Eleni Petraki, Stratos Idreos, Stefan Manegold and Martin L. Kersten. Proceedings of the 10th International Workshop on Data Management on New Hardware (DaMoN), 2014	DaMoN 2015
9	MonetDB/DataCell: Online Analytics in a Streaming Column-Store. Erietta Liarou, Stratos Idreos, Stefan Manegold and Martin L. Kersten. In Proceedings of the Very Large Databases Endowment (PVLDB), 5(12): 1910-1913, 2012	PVLDB 2012
10	NoDB in Action: Adaptive Query Processing on Raw Data. Ioannis Alagiannis, Renata Borovica, Miguel Branco, Stratos Idreos and Anastasia Ailamaki. In Proceedings of the Very Large Databases Endowment (PVLDB), 5(12): 1942-1945, 2012	PVLDB 2012
11	Blink: Not your Father's Database. Ronald Barber, Peter Bendel, Marco Czech, Oliver Draese, Frederick Ho, Namik Hrle, Stratos Idreos, Min-Soo Kim, Oliver Koeth, Jae Gil Lee, Tianchao Tim Lee, Guy Lohman, Konstantinos Morfonios, Keshava Murthy, Lin Qiao, Vijayshankar Raman, Richard Sidle, Knut Stolze. In Proceedings of the 5 th International Workshop on Business Intelligence for Real-Time Enterprises (BIRTE). 2011	BIRTE 2011
12	P2P- DIET: An Extensible P2P Service that Unifies Ad-hoc and Continuous Querying in Super-peer Networks. Stratos Idreos, Manolis Koubarakis and Christos Tryfonopoulos. In Proceedings of the ACM SIGMOD Intern. Conference on Management of Data, 2004	SIGMOD 2004

FUNDING

- 2023 Sloan Research Fellowship. 75K USD
- 2022 Star-Friedman Challenge for Promising Scientific Research. \$84K USD
- 2020 Facebook Research Gift. 150K USD.
- 2020 Tableau, Research Gift, 115K USD.
- 2019 DOE Early Career Award. Data Structure Alchemy. 1M USD.
- 2019 Harvard Data Science Initiative Competitive Research Fund. 69K USD.
- 2019 CISCO. Rapid Training of Deep Neural Network Ensembles. 95K USD.

Stratos Idreos, Curriculum Vitae, January 2024

- 2018 Physical Sciences and Engineering Accelerator Award. 70K USD.
- 2017 Dean's Competitive fund for promising scholarship. 50K USD.
- 2017 Facebook Research Gift. 150K USD.
- 2016, 2018, 2020 NetApp Faculty Fellowship. 150K USD.
- 2016 Facebook Faculty Award. 30K USD.
- 2015-2020 NSF Career Award. Evolutionary Data Systems. US National Foundation. 550K USD.
- 2015-2017 STARnet Research. Near Memory. 300K USD. With David Brooks and Gu –Yeon Wei.
- Industry Gifts: Facebook 10K USD; Google: 1.5K USD.
- 2012 Veni Award. 250K Euros for research on Big Data Exploration. Dutch Organization for Research (NWO)

PATENTS

1. Data Management with Stacked Probabilistic Filters, 2021
2. Optimized Self-Designing Key-Value Storage Engine. 2021
3. Log-Structured Merge-Bush. 2019
4. Diffident Segment Trees In Key-Value Stores: From Point To Range Filters For (Almost) Free. 2019
5. Systems And Methods For Designing Data Structures And Synthesizing Costs. 2018
6. Rapid And Robust Predicate Evaluation. 2018
7. Key-Value Stores With Optimized Merge Policies And Optimized LSM-tree Structures. 2018
8. Data Canopy: Accelerating Exploratory Statistical Analysis. 2017
9. Analytic Database with Optimal Navigable Key-value Store. 2017
10. A Query Management System and Engine Allowing for Efficient Query Execution on Raw Data Files. 2012

SELECTED RECENT TALKS

1. TU Berlin/Berlin Institute for the Foundations of Learning and Data, 2022
2. Perdue University, Database Seminar, 2022
3. UC Berkeley, Database Seminar, 2021
4. German Database conference, Keynote, 2021
5. CISCO Research, 2021
6. Microsoft Research, 2021
7. MIT, ML for Systems class, Invited Lecture, 2021
8. VLDB Workshop on Applied AI for Database Systems and Applications, Keynote, 2021
9. Greek National Student Conference of Electrical and Computer Engineers, Keynote, 2021
10. Facebook Performance Summit IV, 2020
11. ACM Council on Women in Computing, Keynote, 2020
12. Facebook, 2019
13. Google, 2019
14. Microsoft, 2019
15. Oracle, 2019

SERVICE

ORGANIZATION

1. Chair of the ACM SoCC Steering Committee (ACM Symposium on Cloud Computing), 2022-2026
2. Inaugural Editor in Chief. ACM/IMS Journal of Data Science, 2021-2024
3. PC Chair. IEEE International Conference on Data Engineering, 2022
4. PC Chair. ACM SIGMOD International Conference on Management of Data, 2021
5. SIGMOD Executive Committee Member, 2021-2022
6. Advisory Board. ACM SIGMOD International Conference on Management of Data, 2015 - now
7. Reproducibility Chair. ACM SIGMOD International Conference on Management of Data, 2015 –2020
8. Very Large Databases Journal Editorial Board 2020-2026
9. Chair: International Conference on Very Large Databases (VLDB), 2021, Experimental, Analysis and Benchmark Track.
10. PC Chair. North East Database Day 2020
11. PC Chair. Hellenic Data Management Symposium, 2019
12. Workshops Chair, Intern. Conf. on Extending Database Technology (ICDE), 2018
13. Vice Chair, Intern. Conf. on Extending Database Technology (ICDE), 2017
14. PhD workshop Chair, Intern. Conf. on Extending Database Technology (ICDE), 2017
15. Sponsorship Chair. ACM SIGMOD International Conference on Management of Data, 2016 – 2017
16. Chair of Panels. 43th International Conference on Very Large Databases (VLDB), 2017, Munich

Stratos Idreos, Curriculum Vitae, January 2024

17. Chair of Demonstrations Track. 42th International Conference on Very Large Databases (VLDB), 2016
18. Chair. International Workshop on Data Management on New Hardware (DaMoN), 2016/17 with ACM SIGMOD
19. Chair of Tutorials Track. 31th IEEE International Conference on Data Engineering (ICDE), 2015
20. Chair of Demonstrations Track. 17th International Conference on Extending Database Technology (EDBT), 2014

COMMITTEE MEMBER

21. ACM SIGMOD International Conference on Management of Data, 2020, Oregon, USA
22. ACM SIGMOD International Conference on Management of Data, 2019, Amsterdam, The Netherlands
23. ACM SIGMOD International Conference on Management of Data, 2018, Texas, USA
24. ACM SIGMOD International Conference on Management of Data, 2017, North Carolina, USA
25. 8th International Conference on Innovative Data Systems Research (CIDR), 2017, Santa Cruz, CA, USA
26. ACM SIGMOD International Conference on Management of Data, 2016, San Francisco, California, USA
27. IEEE International Conference on Data Engineering (ICDE), 2016, Helsinki, Finland
28. ACM International Conference on Information and Knowledge Management (CIKM), 2015, Melbourne, Australia
29. National Science Foundation, Directorate for Computer & Information Science & Engineering, (NSF CISE), 2015
30. International Journal on Very Large Databases (VLDBJ), 2015
31. ACM Transactions on Database Systems (TODS), 2015
32. 41th International Conference on Very Large Databases (VLDB), 2015, Hawaii, USA
33. 7th International Conference on Innovative Data Systems Research (CIDR), 2015, Monterey, CA, USA
34. 13th USENIX File and Storage Technologies (FAST) Conference, 2015, Santa Clara, CA, USA
35. International Conference on Extending Database Technology (EDBT), 2015, Brussels, Belgium
36. ACM SIGMOD International Conference on Management of Data, 2015, Melbourne, Australia
37. International Workshop on Cloud Data and Platforms (CloudDP), with EuroSys, Bordeaux, France, 2015
38. International Workshop on Exploratory Search in Databases and the Web (ExploreDB), 2015, with ACM SIGMOD
39. International Workshop on Big Data Management on Emerging Hardware (HARDS), with IEEE ICDE, Seoul 2015
40. National Science Foundation, Directorate for Computer & Information Science & Engineering, (NSF CISE), 2014
41. 10th International Workshop on Data Management on New Hardware (DaMoN), 2014, with ACM SIGMOD,
42. ACM International Conference on Information and Knowledge Management (CIKM), 2014, Shanghai, China
43. International Workshop on Self-Managing Database Systems (SMDB), 2014, with IEEE ICDE, Chicago, USA
44. 1st International Workshop on Data analytics in the Cloud (DanaC), 2014, with ACM SIGMOD, Snowbird Utah, USA
45. 30th IEEE International Conference on Data Engineering (ICDE), 2014, Chicago, USA
46. 40th International Conference on Very Large Databases (VLDB), 2014, Hangzhou, China
47. 9th International Workshop on Self-Managing Database Systems (SMDB), 2014, Chicago, USA
48. IEEE International Conference on Big Data (IEEE BigData), 2013, Silicon Valley, CA, USA
49. 33rd ACM SIGMOD International Conference on Management of Data 2013, New York, USA
50. 6th International Conference on Innovative Data Systems Research (CIDR), 2013, Asilomar, California, USA
51. 39th International Conference on Very Large Databases (VLDB), 2013, Trento, Italy
52. 16th International Conference on Extending Database Technology (EDBT), 2013, Genoa, Italy
53. 8th International Workshop on Self Managing Database Systems (SMDB), 2013, Brisbane, Australia
54. 29th British National Conference on Databases (BNCOD), 2013, Oxford, UK
55. 6th International Workshop on Testing Database Systems (DBTest), 2013, New York, USA
56. International Workshop on Linked Science (LISC), 2013, Sydney, Australia
57. 2nd International Conference on Data Management Technologies and Applications (DATA) 2013, Reykjavik, Iceland
58. E.W.R. Steacie Memorial Fellowship, Natural Sciences and Engineering Research Council of Canada (NSERC)'13
59. International Journal of Big Data (IJBD), Science and Engineering Research Support society
60. ACM SIGMOD Jim Gray Doctoral Dissertation Award, 2012
61. 32nd ACM SIGMOD International Conference on Management of Data 2012, Arizona, USA
62. 38th International Conference on Very Large Databases (VLDB) 2012, Istanbul, Turkey
63. 28th IEEE International Conference on Data Engineering (ICDE) 2012, Washington DC, USA
64. 14th International Conference on Extending Database Technology (EDBT) 2011, Uppsala, Sweden
65. 27th IEEE International Conference on Data Engineering (ICDE) 2011, Ph.D. Workshop, Hanover, Germany
66. 6th International Workshop on Self Managing Database Systems (SMDB) 2011, Hanover, Germany
67. 31st ACM SIGMOD International Conference on Management of Data 2011, Repeatability, Athens, Greece
68. 12th International Conference on Extending Database Technology (EDBT) 2009, Saint-Petersburg, Russia
69. IEEE journal on Transactions on Knowledge and Data Engineering (TKDE)