

STRATOS IDREOS

Assistant Professor in Computer Science
Harvard John A. Paulson School of Engineering and Applied Sciences
stratos.seas.harvard.edu stratos@seas.harvard.edu daslab.seas.harvard.edu

Ph.D.: June 2010
Citations > 2800
Hindex:28, i10:56

MAPPING DATA SYSTEMS DNA - SELF-DESIGNING DATA SYSTEMS - HARDWARE/SOFTWARE CO-DESIGN DATA EXPLORATION - RESEARCHER/DEVELOPER PRODUCTIVITY - ACCELERATING DATA SCIENCE

AWARDS



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 Awar. 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 Prize ACM SIGMOD 2018 Undergraduate Research Competition. Winner: T. Lively, C. Mendizabal, L. Schroeder**1st Prize ACM SIGMOD 2017 Undergraduate Research Competition.** Winner: Graham Lustiber**1st Prize ACM SIGMOD 2016 Undergraduate Research Competition.** Winner: Alex Liu**Microsoft PhD Fellowship** 2016. Winner: Lukas Maas.

EMPLOYMENT

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

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, Lukas Maas, Abdul Wasay, Michael Kester, Wilson Qin	PhD	
Konstantinos Zoumpatianos, Niv Dayan, Manos Athanassoulis	Postdoc	
Demi Guo, Drhuv Gupta, Tarik Moon, Stella Pantella, Alex Liu, Wilson Qin, Ore Babarinsa, Yifan Wu	Undergrad	
Data Systems	Spring 2014, Fall 2015, 2016, 2017	http://daslab.seas.harvard.edu/classes/cs165/
Big Data Systems	Fall 2014, Spring 2016, 2017, 2018	http://daslab.seas.harvard.edu/classes/cs265/

PUBLICATIONS

REFEREED CONFERENCES and JOURNALS

- r1. Optimal Bloom Filters and Adaptive Merging for LSM-Trees. Niv Dayan, Manos Athanassoulis, Stratos Idreos. *ACM Transactions on Database Systems*, Dec 2018 **TODS 2018**
- r2. The Periodic Table of Data Structures. Stratos Idreos, Kostas Zoumpatianos, Manos Athanassoulis, Niv Dayan, Brian Hentschel, Michael S. Kester, Demi Guo, Lukas Maas, Wilson Qin, Abdul Wasay, Yiyoun Sun. *IEEE Data Engineering Bull. Sep*, 2018 **IEEE DEBull 2018**
- r3. 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**
- r4. 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**
- r5. 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**
- r6. 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**
- r7. 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 **SIGMOD 2017**
Best of SIGMOD.
- r8. 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**
- r9. 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**
- r10. 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**
- r11. GeckoFTL: Scalable Flash Translation Techniques For Very Large Flash Devices. Niv Dayan, Philippe Bonnet, Stratos Idreos. In *Proc. of the ACM SIGMOD Intern. Conference on Management of Data*, 2016 **SIGMOD 2016**
- r12. 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**
- r13. 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**
- r14. ADS: The Adaptive Data Series Index. Konstantinos Zoumpatianos, Stratos Idreos, Themis Palpanas. *The International Journal on Very Large Databases (VLDBJ)*, 2016 **VLDBJ 2016**
- r15. 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**
- r16. 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**
- r17. 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**
- r18. 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**
- r19. 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**
- r20. 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**
- r21. 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**

- r22. 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**
- r23. 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**
- r24. 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**
- r25. 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**
- r26. 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**
- r27. 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**
- r28. 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**
- r29. 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**
- r30. Business Analytics in (a) Blink. Ronald Barber, Peter Bendel, Marco Czech, Oliver Draese, Frederick Ho, Namik Hrle, 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**
- r31. 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**
- r32. 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**
- r33. 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**
- r34. 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**
- r35. 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
- r36. 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**
- r37. 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**
- r38. 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**
- r39. 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**
- r40. 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**
- r41. 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**
- r42. 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**
- r43. 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**

- r44. 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**
- r45. 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**
- r46. 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**
- r47. 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*
- r48. 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**
- r49. 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

- t1. 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**
- t2. 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**
- t3. 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

- s1. 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**
- s2. Main Memory Adaptive Denormalization. Alex Liu, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2016 **SIGMOD 2016**
- s3. 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**
- s4. 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**
- s5. 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**
- s6. 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**
- s7. 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**
- s8. 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**
- s9. 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**
- s10. 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**
- s11. 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 5th International Workshop on Business Intelligence for Real-Time Enterprises (BIRTE). 2011 **BIRTE 2011**
- s12. 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**

SERVICE

ORGANIZATION

1. Advisory Board. ACM SIGMOD International Conference on Management of Data, 2015 -
2. Reproducibility Chair. ACM SIGMOD International Conference on Management of Data, 2015 -
3. Workshops Chair, Intern. Conf. on Extending Database Technology (ICDE), 2018
4. Vice Chair, Intern. Conf. on Extending Database Technology (ICDE), 2017
5. PhD workshop Chair, Intern. Conf. on Extending Database Technology (ICDE), 2017
6. Sponsorship Chair. ACM SIGMOD International Conference on Management of Data, 2016 - 2017
7. Chair of Panels. 43th International Conference on Very Large Databases (VLDB), 2017, Munich
8. Chair of Demonstrations Track. 42th International Conference on Very Large Databases (VLDB), 2016
9. Chair. International Workshop on Data Management on New Hardware (DaMoN), 2016/17 with ACM SIGMOD
10. Chair of Tutorials Track. 31th IEEE International Conference on Data Engineering (ICDE), 2015
11. Chair of Demonstrations Track. 17th International Conference on Extending Database Technology (EDBT), 2014

COMMITTEE MEMBER

12. ACM SIGMOD International Conference on Management of Data, 2019, Amsterdam, The Netherlands
13. ACM SIGMOD International Conference on Management of Data, 2018, Texas, USA
14. ACM SIGMOD International Conference on Management of Data, 2017, North Carolina, USA
15. 8th International Conference on Innovative Data Systems Research (CIDR), 2017, Santa Cruz, CA, USA
16. ACM SIGMOD International Conference on Management of Data, 2016, San Francisco, California, USA
17. IEEE International Conference on Data Engineering (ICDE), 2016, Helsinki, Finland
18. ACM International Conference on Information and Knowledge Management (CIKM), 2015, Melbourne, Australia
19. National Science Foundation, Directorate for Computer & Information Science & Engineering, (NSF CISE), 2015
20. International Journal on Very Large Databases (VLDBJ), 2015
21. ACM Transactions on Database Systems (TODS), 2015
22. 41th International Conference on Very Large Databases (VLDB), 2015, Hawaii, USA
23. 7th International Conference on Innovative Data Systems Research (CIDR), 2015, Monterey, CA, USA
24. 13th USENIX File and Storage Technologies (FAST) Conference, 2015, Santa Clara, CA, USA
25. International Conference on Extending Database Technology (EDBT), 2015, Brussels, Belgium
26. ACM SIGMOD International Conference on Management of Data, 2015, Melbourne, Australia
27. International Workshop on Cloud Data and Platforms (CloudDP), with EuroSys, Bordeaux, France, 2015
28. International Workshop on Exploratory Search in Databases and the Web (ExploreDB), 2015, with ACM SIGMOD
29. International Workshop on Big Data Management on Emerging Hardware (HARDS), with IEEE ICDE, Seoul 2015
30. National Science Foundation, Directorate for Computer & Information Science & Engineering, (NSF CISE), 2014
31. 10th International Workshop on Data Management on New Hardware (DaMoN), 2014, with ACM SIGMOD,
32. ACM International Conference on Information and Knowledge Management (CIKM), 2014, Shanghai, China
33. International Workshop on Self-Managing Database Systems (SMDB), 2014, with IEEE ICDE, Chicago, USA
34. 1st International Workshop on Data analytics in the Cloud (DanaC), 2014, with ACM SIGMOD, Snowbird Utah, USA
35. 30th IEEE International Conference on Data Engineering (ICDE), 2014, Chicago, USA
36. 40th International Conference on Very Large Databases (VLDB), 2014, Hangzhou, China
37. 9th International Workshop on Self-Managing Database Systems (SMDB), 2014, Chicago, USA
38. IEEE International Conference on Big Data (IEEE BigData), 2013, Silicon Valley, CA, USA
39. 33rd ACM SIGMOD International Conference on Management of Data 2013, New York, USA
40. 6th International Conference on Innovative Data Systems Research (CIDR), 2013, Asilomar, California, USA
41. 39th International Conference on Very Large Databases (VLDB), 2013, Trento, Italy
42. 16th International Conference on Extending Database Technology (EDBT), 2013, Genoa, Italy
43. 8th International Workshop on Self Managing Database Systems (SMDB), 2013, Brisbane, Australia
44. 29th British National Conference on Databases (BNCOD), 2013, Oxford, UK
45. 6th International Workshop on Testing Database Systems (DBTest), 2013, New York, USA
46. International Workshop on Linked Science (LISC), 2013, Sydney, Australia
47. 2nd International Conference on Data Management Technologies and Applications (DATA) 2013, Reykjavik, Iceland
48. E.W.R. Steacie Memorial Fellowship, Natural Sciences and Engineering Research Council of Canada (NSERC)'13
49. International Journal of Big Data (IJBD), Science and Engineering Research Support society
50. ACM SIGMOD Jim Gray Doctoral Dissertation Award, 2012
51. 32nd ACM SIGMOD International Conference on Management of Data 2012, Arizona, USA
52. 38th International Conference on Very Large Databases (VLDB) 2012, Istanbul, Turkey
53. 28th IEEE International Conference on Data Engineering (ICDE) 2012, Washington DC, USA
54. 14th International Conference on Extending Database Technology (EDBT) 2011, Uppsala, Sweden
55. 27th IEEE International Conference on Data Engineering (ICDE) 2011, Ph.D. Workshop, Hanover, Germany
56. 6th International Workshop on Self Managing Database Systems (SMDB) 2011, Hanover, Germany
57. 31st ACM SIGMOD International Conference on Management of Data 2011, Repeatability, Athens, Greece
58. 12th International Conference on Extending Database Technology (EDBT) 2009, Saint-Petersburg, Russia
59. IEEE journal on Transactions on Knowledge and Data Engineering (TKDE)

RECENT TALKS AND LECTURES

- 2018 University of Chicago, New York University, University of Washington, Waterloo University
Vertica, NetApp, Facebook, Google
- 2017 Facebook, Google, Amazon Web Services, NetApp
University of Dresden, Keynote at ExploreDB ACM SIGMOD workshop
- 2016 Hong Kong Chinese University, University of Wisconsin, MIT, North Carolina University
Keynote at the 10th VLDB Workshop on Enabling Real-Time Business Intelligence
- 2015 Brown University, Buffalo University, Paris Descarte University, University of Sydney, Athens University
Keynote at DESWeb workshop on "Data Engineering meets the Semantic Web" in IEEE ICDE

FUNDING

- 2018 Physical Sciences and Engineering Accelerator Award. 70K USD.
- 2017 Dean's Competitive fund for promising scholarship. 50K USD.
- 2017-19 Facebook Research Gift. 150K USD.
- 2016, 2018 NetApp Faculty Fellowship. 100K 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. Key-Value Stores With Optimized Merge Policies And Optimized LSM-tree Structures. 1/2018
2. Data Canopy: Accelerating Exploratory Statistical Analysis. 4/2017
3. Analytic Database with Optimal Navigable Key-value Store. 1/2017
4. A Query Management System and Engine Allowing for Efficient Query Execution on Raw Data Files. 11/2012