

Curriculum Vitae

Stefan Manegold

February 25, 2021

Personalialia

Full name Stefan Manegold

Affiliation Centrum Wiskunde & Informatica (CWI)
Database Architectures (DA)
Science Park 123
1098 XG Amsterdam
The Netherlands

Telephone +31-20-5924212 (office)

E-mail address Stefan.Manegold@CWI.nl
WWW address <http://www.cwi.nl/~manegold/>

Examination

December 2002 *Ph.D in Computer Science*
University of Amsterdam (UvA), The Netherlands
Thesis: *Understanding, Modeling, and Improving Main-Memory Database Performance* [b.02]
Supervisor: Prof. Dr. M. L. Kersten
Committee: Prof. Dr. D. J. DeWitt, Dr. P. van Emde Boas, Prof. Dr. L. O. Hertzberger, Prof. Dr. A. P. J. M. Siebes, Prof. Dr. G. Weikum

September 1994 *M.Sc (Diploma) in Computer Science*
Technical University Clausthal (TUC), Germany
Thesis: *Extending Nested Transactions for the OSCAR Project* [b.01]
Supervisors: Prof. Dr. A. Heuer, Prof. Dr. K. Ecker

Education

1988 – 1994 Technical University Clausthal (TUC), Germany
Major: Informatics
Minor: Automatic Control Technology
Specialization: Relational and object-oriented database systems

1979 – 1988 High School (Gymnasium)
Friedrich-Wilhelm-Schule & Oberstufengymnasium Eschwege, Germany

Professional Experience

since 09/2014 *Professor (0.2 fte) “Data Management”*
Leiden Institute of Advanced Computer Science (LIACS)
Leiden University, Leiden, The Netherlands

since 01/2011 *Head of “Database Architectures” Research Group (full-time, tenure)*
CWI (Centrum Wiskunde & Informatica), Amsterdam, The Netherlands

Leading a successful and productive team of

1 CWI Fellow, 3 tenured researchers, 1 tenure-track researcher, 1 scientific programmer, as well as several post-docs, PhD students, and Master's students.

Projects:

- MonetDB: The Open-source Columnar Database System — Query Processing at Light Speed
- TELEIOS: Virtual Observatory Infrastructure for Earth Observation Data (EU-FP7-ICT)
- COMMIT P19 - *TimeTrails*: Spatiotemporal Data Warehouses for Travellers (NL/FES)
- LEO: Linked Open Earth Observation Data for Precision Farming (EU-FP7-ICT)
- RETHINK big: Roadmap for European Technologies in Hardware and Networking for Big Data (EU-FP7-ICT)
- PROMIMOOC: Process Mining for Multi-objective Online Control (NWO Data Science program “Challenging Big Data”)
- DAMIOSO: Data Mining on High Volume Simulation Output (NWO program “Innovativ Public-Private Partnership in ICT (IPPSI)”); Co-PI
- CIMPLO: Cross-Industry Predictive Maintenance Optimization Platform (NWO/TTW program “Smart Industry”); Co-PI

1997 – 2010

Database researcher (senior/tenure since 2003)

Centrum Wiskunde & Informatica (CWI) in Amsterdam, The Netherlands

Database Group with Prof. Dr. M. L. Kersten

Projects:

- MonetDB: A Novel Main Memory Database Kernel — Query Processing at Light Speed
- CONQUER: Multi-Query Optimization for Data Mining Applications
- MIA: Multimedia Information Analysis
- MultimediaN: Multimedia is the Message
- BRICKS: Basic Research in Informatics for Creating the Knowledge Society

1995 – 1997

Research assistant

Humboldt-University Berlin, Germany, Institute for Informatics

Database Group with Prof. J.-C. Freytag, Ph.D.

Project:

- Query Evaluation and Optimization in Parallel Databases
(funded by the German Research Foundation)

1994 – 1995

Compulsory Military Service

Computing Center of the German Air Force

System administration on Siemens mainframes (H90 & 7.580i) running BS 2000

1991 – 1994

Student-assistant

Technical University Clausthal, Germany, Institute for Informatics

Database Group with Prof. Dr. A. Heuer and

Data-management Group with Prof. Dr. D.P.F. Möller

Projects:

- Design and implementation of a data dictionary and a data definition language for the object-oriented database system OSCAR
- Design of a database schema for a medical information system
- Design and implementation of a data dictionary (meta-schema) for an object-oriented database system on top of ObjectStore

1988 – 1992

Free-lance programmer

Design, implementation, and administration of several dBase applications for a metalworking company

Awards

2020

2020 ACM SIGMOD Contributions Award

(<https://sigmod.org/2020-sigmod-contributions-award/>)

2011

VLDB 2011 Challenges & Visions Track Best Paper Award [d.11, e.32]

2009

VLDB 2009 10-year Best Paper Award [d.08, e.27, e.07, h.14]

Projects

2017 - 2021	CIMPLO: Cross-Industry Predictive Maintenance Optimization Platform (NWO/TTW program “Smart Industry”; Co-PI)
2016 - 2020	DAMIOSO: Data Mining on High Volume Simulation Output (NWO program “Innovativ Public-Private Partnership in ICT (IPPSI)”; Co-PI)
2014 - 2018	PROMIMOOC: Process Mining for Multi-objective Online Control (NWO Data Science program “Challenging Big Data”; Co-PI)
2014 - 2016	RETHINK big: Roadmap for European Technologies in Hardware and Networking for Big Data (EU-FP7-ICT CSA; Co-PI, working group leader)
2013 - 2015	LEO: Linked Open Earth Observation Data for Precision Farming (EU-FP7-ICT STREP; Co-PI, Quality Assurance Coordinator)
2011 - 2017	COMMIT P19 - <i>TimeTrails</i> : Spatiotemporal Data Warehouses for Travellers (NL/FES; work package leader)
2010 - 2013	TELEIOS: Virtual Observatory Infrastructure for Earth Observation Data (EU-FP7-ICT STREP; Quality Assurance Coordinator)

PhD (co-)advisorship

2018-2022	Matheus Agio Nerone	(CWI & Leiden University)
2017-2021	Pedro Thiago Timbo Holanda	(CWI & Leiden University)
2016-2020	Mark Raasveldt	(CWI & Leiden University)
	- <i>Winner Student Research Competition @ ACM SIGMOD 2018 (Houston, TX, USA)</i>	
2011-2018	Minh Duc Pham	(CWI & Vrije Universiteit Amsterdam)
2011-2017	Yağız Kargin	(CWI & University of Amsterdam)
2011-2017	Eleni Petraki	(CWI & University of Amsterdam)
2011-2016	Mrunal Gawade	(CWI & University of Amsterdam)
2011-2015	Thibault Sellam	(CWI & University of Amsterdam)
2010-2014	Holger Pirk	(CWI & University of Amsterdam)
2008-2013	Romulo Antonio Pereira Gonçalves	(CWI & University of Amsterdam)
2006-2012	Erietta Liarou	(CWI & University of Amsterdam)
2005-2010	Stratos Idreos	(CWI & University of Amsterdam)
	- <i>ACM SIGMOD Jim Gray Doctoral Dissertation Award (2011)</i>	
	- <i>ERCIM Cor Baayen Award (2011)</i>	
2004-2009	Fabian Groffen	(CWI & University of Amsterdam)

Professional Activities & Community Services

2021	<ul style="list-style-type: none">- Co-Founder & Steering Committee member of International Workshop Series on Data Management on New Hardware (DaMoN; Sponsored by and colocated with ACM SIGMOD/PODS; since 2005)- Associate Editor of PVLDB Vol. 14 (VLDB 2021)- Vice Chair (Track Chair) of ICDE 2021- Tutorial Chair of EDBT 2021- Program Committee member of BTW 2021- Co-Chair of HardBD & Active 2021- External PhD Examiner Clemens Lutz, Technical University Berlin, Berlin, Germany- External PhD Examiner Dennis Marten, University of Rostock, Rostock, Germany
2020	<ul style="list-style-type: none">- Local Arrangement Chair of CIDR 2020- Program Committee member of VLDB 2020 PhD workshop- Program Committee member of HardBD 2020- External PhD Examiner Oksana Dolmatova, University of Zürich, Zürich, Switzerland- External PhD Examiner Patrick Damme, Technical University Dresden, Dresden, Germany- PhD Examiner Armon Toubman, Leiden University, Leiden, The Netherlands

- 2019
- General Chair of ACM SIGMOD/PODS 2019
 - Member of PVLDB Review Board (Research Track VLDB 2019)
 - Demonstration Program Committee member of EDBT 2019
 - Scientific Program Committee member of BTW 2019
 - Demonstration Program Committee member of BTW 2019
 - Program Committee member of HardBD 2019
 - External PhD Examiner David Broneske, Otto von Guericke University Magdeburg, Magdeburg, Germany
 - PhD Examiner Christian Fuchs, Leiden University, Leiden, The Netherlands
- 2018
- Lightning Talks Track Chair of ICDE 2018
 - Research Track Program Committee member of EDBT 2018
 - Demonstration Program Committee member of SIGMOD 2018
 - Program Committee member of DBTest 2018
 - Program Committee member of ADMS 2018
 - Member of Advisory Committee for TNO program “Making Sense of Big Data”
 - External PhD Examiner Steffen Zeuch, Humboldt University Berlin, Berlin, Germany
 - PhD Examiner Bas van Stein, Leiden University, Leiden, The Netherlands
 - PhD Examiner Hao Wang, Leiden University, Leiden, The Netherlands
- 2017
- Research Track Associate Editor of PVLDB/VLDB 2017
 - Member of PVLDB Review Board
 - Program Committee member of SIGMOD 2017
 - Demonstration Program Committee member of EDBT 2017
 - Demonstration Program Committee member of BTW 2017
 - Program Committee member of ADMS 2017
 - Member of Advisory Committee for TNO program “Making Sense of Big Data”
 - External PhD Examiner Ismail Oukid, Technical Dresden, Dresden, Germany
 - PhD Examiner Yuanhao Guo, Leiden University, Leiden, The Netherlands
- 2016
- Demonstration Program Committee Chair of ICDE 2016
 - Member of PVLDB Review Board
 - Program Committee member of VLDB 2016
 - Program Committee member of ADMS 2016
 - Member of Advisory Committee for TNO program “Making Sense of Big Data”
 - Reviewer for FP-7 project AXLE for the European Commission
- 2015
- Member of PVLDB Review Board
 - Associate Editor “Novel DB Architectures” of PVLDB/VLDB 2015
 - Program Committee member of VLDB 2015 demonstrations track
 - Program Committee member of VLDB 2015 PhD workshop
 - Program Committee member of ICDE 2015
 - Program Committee member of EDBT 2015
 - Program Committee member of IC2E 2015
 - Program Committee member of ADMS 2015
 - Demonstration Program Committee member of BTW 2015
 - PhD Examiner Enrique Larios Vargas, Leiden University, Leiden, The Netherlands
 - External PhD Examiner Florian Funke, Technical University Munich, Munich, Germany
 - Reviewer for FP-7 project AXLE for the European Commission
- 2014
- Member of PVLDB Review Board
 - Program Committee member of VLDB 2014
 - Program Committee member of SIGMOD 2014
 - Program Committee member of ICDE 2014
 - Program Committee member of EDBT 2014
 - Program Committee member of ADMS 2014
 - Program Committee member of IMDM 2014
 - Program Committee member of Terra Cognita 2014

- Member of ICDE 2014 Panel “Main-Memory Database Systems”
 - External PhD Examiner Cagri Balkesen, ETH Zurich, Zurich, Switzerland
 - External PhD Examiner Kostis Kyzirakos, National and Kapodistrian University of Athens, Athens, Greece
 - Reviewer for FP-7 project AXLE for the European Commission
- 2013
- Member of PVLDB Review Board
 - Associate Editor “Experiments and Analysis Track” of PVLDB/VLDB 2013
 - Program Committee member of BTW 2013 (Scientific Program Track)
 - Demonstration Program Committee member of BTW 2013
 - Program Committee member of EDBT 2013
 - Program Committee member of IC2E 2013
 - Program Committee member of SMDB 2013
 - Program Committee member of BigData 2013
 - Program Committee member of ADMS 2013
 - Program Committee member of IMDM 2013
- 2012
- Member of ADMS 2012 Panel “Impact of the convergence of three worlds: High-Performance Computing, Databases, and Analytics”
 - Guest editor for ECRIM News no. 89 special theme on “Big Data”
 - Program Committee member of VLDB 2012
 - Program Committee member of ADMS 2012
 - Program Committee member of ICDE 2012 (Track “Cloud, data warehousing, and large data”)
 - Program Committee member of SMDB 2012
 - Demonstration Program Committee member of EDBT 2012
 - Program Committee member of Terra Cognita 2012
 - Co-Organizer of EDBT 2012 Tutorial “Adaptive Indexing in Modern Databases”
- 2011
- Member of ADMS 2011 Panel “In the multi-core age, how do larger, faster, cheaper and more responsive memory sub-systems affect data management?”
 - Advisor for Repeatability & Workability Evaluation of SIGMOD 2011
 - Program Committee member of Terra Cognita 2011
 - Program Committee member of DBTest 2011
 - External PhD Examiner Gerard Marks, Dublin City University, Dublin, Ireland
 - External PhD Examiner Syed Saif ur Rahman, Otto-von-Guericke University, Magdeburg, Germany
- 2010
- Repeatability & Workability Evaluation Co-Chair of SIGMOD 2010
 - Program Committee member of SIGMOD 2010
 - Program Committee member of ICDE 2010 (Track “Systems, Experiments, Applications”)
 - Program Committee member of EDBT 2010
- 2009
- Repeatability & Workability Evaluation Co-Chair of SIGMOD 2009
 - Demonstration Program Committee Chair of EDBT 2009
 - Co-Organizer of EDBT 2009 Tutorial “Performance Evaluation in Database Research: Principles and Experience”
 - Program Committee member of VLDB 2009 E&A Track
 - Program Committee member of ICDE 2009 (Track “Transaction and workflow management, database architectures”)
 - Program Committee member of SMDB 2009
 - Program Committee member of BTW 2009
- 2008
- Co-Organizer of ICDE 2008 Seminar “Performance Evaluation in Database Research: Principles and Experience”
 - Experimental Repeatability Committee member of SIGMOD 2008
 - Program Committee member of EDBT 2008
 - Demonstration Program Committee member of ICDE 2008
 - Program Committee member of VLDB 2008 PhD Workshop

- 2007
 - Co-Organizer of VLDB 2007 Panel “Performance Evaluation and Experimental Assessment — Conscience or Curse of Database Research?”
 - Co-Organizer of the Second International Workshop on Performance and Evaluation of Data Management Systems (ExpDB 2007; Sponsored by and colocated with ACM SIGMOD/PODS 2007)
 - Program Committee member of DASFAA 2007
 - Program Committee member of DEXA 2007
 - Program Committee member of BTW 2007

- 2006
 - Co-Organizer of the Second International Workshop on Data Management on New Hardware (DaMoN 2006; Sponsored by and colocated with ACM SIGMOD/PODS 2006)
 - Demonstration Program Committee member of VLDB 2006
 - Program Committee member of EDBT 2006 PhD Workshop
 - Program Committee member of QLQP 2006 (EDBT Workshop)
 - Program Committee member of DASFAA 2006

- 2005
 - Co-Organizer of the First International Workshop on Data Management on New Hardware (DaMoN 2005; Sponsored by and colocated with ACM SIGMOD/PODS 2005)
 - Program Committee member of VLDB 2005 (Core Database Technology)

- 2004
 - Program Committee member of SAC 2004 (Database track)

- since 2006
 - Teaching course “Advanced Database Techniques” at the University of Amsterdam

- 2004 – 2009
 - Teaching course “Database Techniques” at the University of Amsterdam

- since 2004
 - Supervising Masters Students & PhD Students

- since 1996
 - Reviewer for various journals, conferences and workshops, a.o.:
EDBT, FGCS, ICDE, Information Systems, PODS, PVLDB, SAC, SIGMOD, TKDE, TODS, TOIS, VLDB, VLDB Journal

Research Interests

- Exploratory Big Data Analytics
- Column-store Database Architectures & Technologies
- Data Management for Sciences
- Analytical Database Systems
- Parallel & Distributed Information Systems
- Cost Models for Database Query Optimization
- XML storage & XQuery processing using RDBMS technology
- Hardware-aware Database Technology
- Processor Architectures & Hierarchical Cache Memories
- Optimization, Scalability, Performance, & Benchmarking
- Multi-Architecture/Cross-Platform Software-Development & -Testing

Technical Skills

General skills in computer science. Specific expertise in:

- Software development, release management, performance analysis & monitoring, regression testing, and quality assurance for open-source database management system MonetDB (<http://www.monetdb.org/>)
- Computing platforms: several UNIX variants, with a focus on Linux, and some Windows experience.

- Cross-platform software development: Linux (on AMD, Intel & StrongARM), Solaris, IRIX, AIX, MacOS X, CygWin, Windows
- System administration: cluster of Linux PCs, Siemens mainframes (H90 & 7.580i) running BS 2000
- Programming languages: fluent in C, C++, Pascal, Basic, SQL, dBase, Python, (ba)sh; some experience in Perl, prolog, assembly, Java
- Parallel programming experience in C/C++ using PVM and MPI on a cluster of Sun workstations and on a CRAY T3E

Languages

German (native), English (fluent), Dutch (fluent), French (basic)

Publications

a. Proceedings (edited)

- [a.05] P. A. Boncz, **S. Manegold**, A. Ailamaki, A. Deshpande, and T. Kraska, editors. *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, Amsterdam, The Netherlands, June 2019.
- [a.04] M. L. Kersten, B. Novikov, J. Teubner, V. Polutin, and **S. Manegold**, editors. *12th International Conference on Extending Database Technology (EDBT 2009)*, Saint Petersburg, Russia, March 2009.
- [a.03] P. Bonnet and **S. Manegold**, editors. *Second International ACM/SIGMOD Workshop on Performance and Evaluation of Data Management Systems (ExpDB 2007)*, Beijing, China, June 2007.
- [a.02] A. Ailamaki, P. A. Boncz, and **S. Manegold**, editors. *Second International ACM/SIGMOD Workshop on Data Management on New Hardware (DaMoN 2006)*, Chicago, IL, USA, June 2006.
- [a.01] A. Ailamaki, P. A. Boncz, and **S. Manegold**, editors. *First International ACM/SIGMOD Workshop on Data Management on New Hardware (DaMoN 2005)*, Baltimore, MD, USA, June 2005.

b. Theses

- [b.02] **S. Manegold**. *Understanding, Modeling, and Improving Main-Memory Database Performance*. PhD thesis, Universiteit van Amsterdam, Amsterdam, The Netherlands, December 2002.
- [b.01] **S. Manegold**. Extending Nested Transactions for the OSCAR-Project. Master's thesis, Technische Universität Clausthal, Institut für Informatik, Clausthal-Zellerfeld, Germany, September 1994. *In German*.

c. Book Chapters

- [c.09] **S. Manegold**. Cost Estimation. In L. Liu and M. T. Özsu, editors, *Encyclopedia of Database Systems*, pages 506–511. Springer-Verlag, Heidelberg, Germany, 2nd edition, October 2018.
- [c.08] **S. Manegold**. Memory Hierarchy. In L. Liu and M. T. Özsu, editors, *Encyclopedia of Database Systems*, pages 1707–1713. Springer-Verlag, Heidelberg, Germany, 2nd edition, October 2018.
- [c.07] **S. Manegold**. Locality. In L. Liu and M. T. Özsu, editors, *Encyclopedia of Database Systems*, pages 1713–1714. Springer-Verlag, Heidelberg, Germany, 2nd edition, October 2018.
- [c.06] **S. Manegold**. Cost Estimation. In L. Liu and M. T. Özsu, editors, *Encyclopedia of Database Systems*, pages 506–511. Springer-Verlag, Heidelberg, Germany, 1st edition, October 2009.
- [c.05] **S. Manegold**. Memory Hierarchy. In L. Liu and M. T. Özsu, editors, *Encyclopedia of Database Systems*, pages 1707–1713. Springer-Verlag, Heidelberg, Germany, 1st edition, October 2009.
- [c.04] **S. Manegold**. Locality. In L. Liu and M. T. Özsu, editors, *Encyclopedia of Database Systems*, pages 1713–1714. Springer-Verlag, Heidelberg, Germany, 1st edition, October 2009.

- [c.03] A. R. Schmidt, **S. Manegold**, and M. L. Kersten. Storing XML Documents in Databases. In L. C. Rivero, J. H. Doorn, and V. E. Ferragline, editors, *Encyclopedia of Database Technologies and Applications*, pages 658–664. Idea Group Publishing, August 2005.
- [c.02] A. R. Schmidt, F. Waas, **S. Manegold**, and M. L. Kersten. A Look Back on the XML Benchmark Project. In H. M. Blanken, T. Grabs, H.-J. Schek, R. Schenkel, and G. Weikum, editors, *Intelligent Search on XML*, volume 2818 of *Lecture Notes in Computer Science*, chapter 18, pages 263–278. Springer-Verlag, Berlin, New York, etc., August 2003.
- [c.01] **S. Manegold** and F. Waas. Integrating I/O processing and Transparent Parallelism — Toward Comprehensive Query Execution in Parallel Database Systems. In A. Dogaç, M. T. Özsu, and Ö. Ulusoy, editors, *Current Trends in Data Management Technology*, chapter 8, pages 130–152. Idea Group Publishing, Hershey, PA, USA, January 1999.

d. Journals

- [d.30] P. Holanda, **S. Manegold**, H. Mühleisen, and M. Raasveldt. Progressive Indexes: Indexing for Interactive Data Analysis. *Proceedings of the VLDB Endowment (PVLDB)*, 12(13):2366–2378, September 2019. (Same as [e.63].)
- [d.29] K. Kyzirakos, D. Savva, I. Vlachopoulos, A. Vasileiou, N. Karalis, M. Koubarakis, and **S. Manegold**. GeoTriples: Transforming geospatial data into RDF graphs using R2RML and RML mappings. *Journal of Web Semantics*, September 2018.
- [d.28] M. Koubarakis, K. Kyzirakos, C. Nikolaou, G. Garbis, K. Bereta, R. Dogani, S. Giannakopoulou, P. Smeros, D. Savva, G. Stamoulis, G. Vlachopoulos, **S. Manegold**, C. Kontoes, T. Herekakis, I. Papoutsis, and D. Michail. Managing Big, Linked, and Open Earth-Observation Data: Using the TELEIOS/LEO software stack. *IEEE Geoscience and Remote Sensing Magazine*, 4(3):23–37, September 2016.
- [d.27] R. Cijvat, **S. Manegold**, M. L. Kersten, G. W. Klau, A. Schönhuth, T. Marschall, and Y. Zhang. Genome Sequence Analysis with MonetDB: A Case Study on Ebola Virus Diversity. *Datenbank-Spektrum*, 15(3):185–191, November 2015. *Best Workshop Papers of BTW 2015*.
- [d.26] **S. Manegold**. Front Matter: Letter from the Guest Editors. *Proceedings of the VLDB Endowment (PVLDB)*, 8(7):8, March 2015.
- [d.25] G. Graefe, F. Halim, S. Idreos, H. Kuno, **S. Manegold**, and B. Seeger. Transactional Support for Adaptive Indexing. *The VLDB Journal*, 23(2):303–328, April 2014. *Special Issue: Best papers of VLDB 2012*.
- [d.24] K. Kyzirakos, **S. Manegold**, C. Nikolaou, and M. Koubarakis. Revolutionary Database Technology for Data Intensive Research. *ERCIM News*, 2014(96), January 2014. *Special Theme on Linked Open Data*.
- [d.23] Y. Kargin, M. Ivanova, Y. Zhang, **S. Manegold**, and M. L. Kersten. Lazy ETL in Action: ETL Technology Dates Scientific Data. *Proceedings of the VLDB Endowment (PVLDB)*, 12(9):1286–1289, August 2013. *Demo*. (Same as [e.50].)
- [d.22] M. Heimel, M. Saecker, H. Pirk, **S. Manegold**, and V. Markl. Hardware-Oblivious Parallelism for In-Memory Column-Stores. *Proceedings of the VLDB Endowment (PVLDB)*, 6(9):709–720, July 2013. (Same as [e.49].)
- [d.21] S. Amer-Yahia and **S. Manegold**. Front Matter: Letter from the Guest Editors. *Proceedings of the VLDB Endowment (PVLDB)*, 6(8):10, June 2013.
- [d.20] M. Ivanova, M. L. Kersten, **S. Manegold**, and Y. Kargin. Data Vaults: Database Technology for Scientific File Repositories. *Computing in Science and Engineering (CiSE)*, 15(3):32–42, May 2013.
- [d.19] E. Liarou, S. Idreos, **S. Manegold**, and M. L. Kersten. MonetDB/DataCell: Online Analytics in a Streaming Column-Store. *Proceedings of the VLDB Endowment (PVLDB)*, 5(12):1910–1913, August 2012. *Demo*. (Same as [e.40].)
- [d.18] M. Koubarakis, K. Kyzirakos, M. Karpathiotakis, C. Nikolaou, S. Vassos, G. Garbis, M. Sioutis, K. Bereta, D. Michail, C. Kontoes, I. Papoutsis, T. Herekakis, **S. Manegold**, M. L. Kersten, M. Ivanova, H. Pirk, Y. Zhang, M. Datcu, G. Schwarz, C. Dumitru, D. Espinoza-Molina, K. Molch, U. Di Giammatteo, M. Sagona, S. Perelli, T. Reitz, E. Klien, and R. Gregor. TELEIOS: A Database-Powered Virtual Earth Observatory. *Proceedings of the VLDB Endowment (PVLDB)*, 5(12):2010–2013, August 2012. *Demo*. (Same as [e.39].)
- [d.17] C. Thanos, **S. Manegold**, and M. L. Kersten. Big Data - Introduction to the special theme. *ERCIM News*, 2012(89):10–11, April 2012. *Special Theme on Big Data*.

- [d.16] M. L. Kersten and **S. Manegold**. Revolutionary Database Technology for Data Intensive Research. *ERCIM News*, 2012(89):16, April 2012. *Special Theme on Big Data*.
- [d.15] S. Idreos, F. Groffen, N. Nes, **S. Manegold**, K. S. Mullender, and M. L. Kersten. MonetDB: Two Decades of Research on Column-oriented Database Architectures. *IEEE Data Engineering Bulletin*, 35(1):40–45, March 2012. *Special Issue on Column Store Systems*.
- [d.14] G. Graefe, F. Halim, S. Idreos, H. Kuno, and **S. Manegold**. Concurrency Control for Adaptive Indexing. *Proceedings of the VLDB Endowment (PVLDB)*, 5(7):656–667, March 2012. (Same as [e.38].)
- [d.13] M. L. Kersten, **S. Manegold**, and K. S. Mullender. The Database Architectures Research Group at CWI. *ACM SIGMOD Record*, 40(4):39–44, December 2011.
- [d.12] P. Bonnet, **S. Manegold**, M. Bjorling, W. Cao, J. Gonzales, J. Granados, N. Hall, S. Idreos, M. Ivanova, R. Johnson, D. Koop, T. Kraska, R. Müller, D. Olteanu, P. Papotti, C. Reilly, D. Tsirogiannis, C. Yu, J. Freire, and D. Shasha. Repeatability & Workability Evaluation of SIGMOD 2011. *ACM SIGMOD Record*, 40(2):45–48, June 2011.
- [d.11] M. L. Kersten, S. Idreos, **S. Manegold**, and E. Liarou. The Researcher’s Guide to the Data Deluge: Querying a Scientific Database in Just a Few Seconds. *Proceedings of the VLDB Endowment (PVLDB)*, 4(12):1474–1477, August 2011. *Challenges & Visions Track Best Paper Award*. (Same as [e.32].)
- [d.10] S. Idreos, **S. Manegold**, H. Kuno, and G. Graefe. Merging What’s Cracked, Cracking What’s Merged: Adaptive Indexing in Main-Memory Column-Stores. *Proceedings of the VLDB Endowment (PVLDB)*, 4(9):585–597, June 2011. (Same as [e.31].)
- [d.09] **S. Manegold**, I. Manolescu, L. Afanasiev, J. Feng, G. Gou, M. Hadjieleftheriou, S. Harizopoulos, P. Kalnis, K. Karanasos, D. Laurent, M. Lupu, N. Onose, C. Re, V. Sans, P. Senellart, T. Wu, and D. Shasha. Repeatability & Workability Evaluation of SIGMOD 2009. *ACM SIGMOD Record*, 38(3):40–43, September 2009.
- [d.08] **S. Manegold**, M. L. Kersten, and P. A. Boncz. Database Architecture Evolution: Mammals Flourished long before Dinosaurs became Extinct. *Proceedings of the VLDB Endowment (PVLDB)*, 2(2):1648–1653, August 2009. *10-year Best Paper Award paper for [e.07]*. (Same as [e.27].)
- [d.07] P. A. Boncz, M. L. Kersten, and **S. Manegold**. Breaking the Memory Wall in MonetDB. *Communications of the ACM (CACM)*, 51(12):77–85, December 2008. In “Research Highlights”.
- [d.06] L. Sidirourgos, R. Gonçalves, M. L. Kersten, N. Nes, and **S. Manegold**. Column-Store Support for RDF Data Management: not all swans are white. *Proceedings of the VLDB Endowment (PVLDB)*, 1(2):1553–1563, August 2008. (Same as [e.23].)
- [d.05] I. Manolescu, L. Afanasiev, A. Arion, J.-P. Dittrich, **S. Manegold**, N. Polyzotis, K. Schnaitter, P. Senellart, S. Zoupanos, and D. Shasha. The repeatability experiment of SIGMOD 2008. *ACM SIGMOD Record*, 37(1):39–45, March 2008.
- [d.04] **S. Manegold**. An Empirical Evaluation of XQuery Processors. *Information Systems*, 33(2):203–220, April 2008. *Updated and extended version of [e.19]*.
- [d.03] A. Ailamaki, P. A. Boncz, and **S. Manegold**. Report on the Second International Workshop on Data Management on Modern Hardware (DaMoN’06). *ACM SIGMOD Record*, 35(4):46–47, December 2006.
- [d.02] **S. Manegold**, P. A. Boncz, and M. L. Kersten. Optimizing Main-Memory Join On Modern Hardware. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 14(4):709–730, July 2002.
- [d.01] **S. Manegold**, P. A. Boncz, and M. L. Kersten. Optimizing Database Architecture for the New Bottleneck: Memory Access. *The VLDB Journal*, 9(3):231–246, December 2000.

e. Conferences & Workshops

- [e.66] M. Agio Nerone, P. Holanda, E. C. de Almeida, and **S. Manegold**. Multidimensional Adaptive & Progressive Indexes. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE)*, Chania, Crete, Greece, April 2021.
- [e.65] P. Holanda and **S. Manegold**. Progressive Mergesort: Merging Batches of Appends into Progressive Indexes. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, page 6, Nicosia, Cyprus, March 2021.

- [e.64] M. Ding, S. Chen, N. Makrynioti, and **S. Manegold**. Progressive Join Algorithms Considering User Preference. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*, online, January 2021.
- [e.63] P. Holanda, **S. Manegold**, H. Mühleisen, and M. Raasveldt. Progressive Indexes: Indexing for Interactive Data Analysis. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Tokyo, Japan, August 2020. (Same as [d.30].)
- [e.62] M. Nerone, **S. Manegold**, and H. Pirk. Efficient Hard Real-Time Transaction Processing — Unachievable in Software. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*, Amsterdam, The Netherlands, January 2020.
- [e.61] M. Raasveldt, P. Holanda, and **S. Manegold**. devUDF: Increasing UDF development efficiency through IDE Integration. It works like a PyCharm! In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, pages 558–561, Lisbon, Portugal, March 2019. *Demo*.
- [e.60] M. L. Kersten, **S. Manegold**, Y. Zhang, and P. Kuoutsourakis. SQLPEL: A database performance platform. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*, Asilomar, CA, USA, January 2019.
- [e.59] P. Holanda, M. Nerone, E. Cunha de Almeida, and **S. Manegold**. Cracking KD-Tree: The First Multidimensional Adaptive Indexing. In *Proceedings of the International Conference on Data Science, Technology and Applications (DATA)*, pages 393–399, Porto, Portugal, July 2018.
- [e.58] M. Raasveldt, P. Holanda, H. Mühleisen, and **S. Manegold**. Deep Integration of Machine Learning Into Column Stores. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, pages 473–476, Vienna, Austria, March 2018.
- [e.57] E. Petraki, S. Idreos, and **S. Manegold**. Holistic Indexing in Main-memory Column-stores. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, Melbourne, Australia, June 2015.
- [e.56] Y. Kargin, M. L. Kersten, **S. Manegold**, and H. Pirk. The DBMS – your Big Data Sommelier. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE)*, Seoul, Korea, April 2015.
- [e.55] R. Cijvat, **S. Manegold**, M. L. Kersten, G. W. Klau, A. Schönhuth, T. Marschall, and Y. Zhang. Genome Sequence Analysis with MonetDB: A Case Study on Ebola Virus Diversity. In *Proceedings of the GI-Fachtagung für Datenbanksysteme in Business, Technologie und Web (BTW) – Workshopband*, volume 242 of *Lecture Notes in Informatics*, pages 143–150, Hamburg, Germany, March 2015. Gesellschaft für Informatik. Joint Workshop on Data Management for Science (DMS).
- [e.54] H. Mühleisen, M. L. Kersten, and **S. Manegold**. Capturing the Laws of (Data) Nature. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*, Asilomar, CA, USA, January 2015.
- [e.53] M. Koubarakis, P. Smeros, C. Nikolaou, G. Garbis, K. Bereta, S. Giannakopoulou, K. Dogani, M. Karpathiotaki, I. Vlachopoulos, D. Savva, G. Stamoulis, K. Kyzirakos, **S. Manegold**, B. Valentin, H. Bach, F. Niggemann, P. Klug, W. Angermair, and S. Burgstaller. Linked Open Data in the Earth Observation Domain: the Vision of Project LEO. In *Conference on Big Data from Space (BiDS)*, Frascati, Italy, November 2014.
- [e.52] H. Pirk, E. Petraki, S. Idreos, **S. Manegold**, and M. L. Kersten. Database Cracking: Fancy Scan, not Poor Man’s Sort! In *Proceedings of the International Workshop on Data Management on New Hardware (DaMoN)*, Snowbird, Utah, USA, June 2014.
- [e.51] H. Pirk, **S. Manegold**, and M. L. Kersten. Waste Not... Efficient Co-Processing of Relational Data. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE)*, Chicago, IL, USA, April 2014.
- [e.50] Y. Kargin, M. Ivanova, Y. Zhang, **S. Manegold**, and M. L. Kersten. Lazy ETL in Action: ETL Technology Dates Scientific Data. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Riva del Garda, Italy, August 2013. *Demo*. (Same as [d.23].)
- [e.49] M. Heimel, M. Saecker, H. Pirk, **S. Manegold**, and V. Markl. Hardware-Oblivious Parallelism for In-Memory Column-Stores. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Riva del Garda, Italy, August 2013. (Same as [d.22].)
- [e.48] M. Ivanova, Y. Kargin, M. L. Kersten, **S. Manegold**, Y. Zhang, M. Datcu, and D. Espinoza-Molina. Data Vaults: A Database Welcome to Scientific File Repositories. In *Proceedings of the International Conference on Scientific and Statistical Database Management (SSDBM)*, page 48, Baltimore, MD, USA, July 2013. *Demo*.

- [e.47] Y. Zhang, M. L. Kersten, and **S. Manegold**. SciQL: Array Data Processing Inside An RDBMS. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, New York, NY, USA, June 2013. *Demo*.
- [e.46] H. Pirk, F. Funke, M. Grund, T. Neumann, U. Leser, **S. Manegold**, A. Kemper, and M. L. Kersten. CPU and Cache Efficient Management of Memory-Resident Databases. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE)*, pages 14–25, Brisbane, Australia, April 2013.
- [e.45] E. Liarou, S. Idreos, **S. Manegold**, and M. L. Kersten. Enhanced Stream Processing in a DBMS Kernel. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, Genua, Italy, March 2013.
- [e.44] M. Koubarakis, C. Kontoes, **S. Manegold**, M. Karpathiotakis, K. Kyzirakos, K. Bereta, G. Garbis, C. Nikolaou, D. Michail, I. Papoutsis, T. Herekakis, M. Ivanova, Y. Zhang, H. Pirk, M. L. Kersten, K. Dogani, S. Gianakopoulou, and P. Smeros. Real-Time Wildfire Monitoring Using Scientific Database and Linked Data Technologies. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, Genua, Italy, March 2013.
- [e.43] M. Koubarakis, M. Karpathiotakis, K. Kyzirakos, C. Nikolaou, S. Vassos, G. Garbis, M. Sioutis, K. Bereta, **S. Manegold**, M. L. Kersten, M. Ivanova, H. Pirk, Y. Zhang, C. Kontoes, I. Papoutsis, T. Herekakis, D. Michail, M. Datcu, G. Schwarz, O. C. Dumitru, D. E. Molina, K. Molch, U. Di Giammatteo, M. Sagona, S. Perelli, E. Klien, T. Reitz, and R. Gregor. Building Virtual Earth Observatories using Ontologies and Linked Geospatial Data. In *Proceedings of the International Conference on Web Reasoning and Rule Systems (RR)*, Vienna, Austria, September 2012.
- [e.42] Y. Kargin, H. Pirk, M. Ivanova, **S. Manegold**, and M. L. Kersten. Instant-On Scientific Data Warehouses: Lazy ETL for Data-Intensive Research. In *Proceedings of the International Workshop on Business Intelligence for the Real Time Enterprise (BIRTE)*, Istanbul, Turkey, August 2012.
- [e.41] K. Bösche, T. Sellam, H. Pirk, R. Beier, P. Mieth, and **S. Manegold**. Scalable Generation of Synthetic GPS Traces with Real-life Data Characteristics. In *Proceedings of the TPC Technology Conference on Performance Evaluation & Benchmarking (TPC TC)*, Istanbul, Turkey, August 2012.
- [e.40] E. Liarou, S. Idreos, **S. Manegold**, and M. L. Kersten. MonetDB/DataCell: Online Analytics in a Streaming Column-Store. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Istanbul, Turkey, August 2012. *Demo*. (Same as [d.19].)
- [e.39] M. Koubarakis, K. Kyzirakos, M. Karpathiotakis, C. Nikolaou, S. Vassos, G. Garbis, M. Sioutis, K. Bereta, D. Michail, C. Kontoes, I. Papoutsis, T. Herekakis, **S. Manegold**, M. L. Kersten, M. Ivanova, H. Pirk, Y. Zhang, M. Datcu, G. Schwarz, C. Dumitru, D. Espinoza-Molina, K. Molch, U. Di Giammatteo, M. Sagona, S. Perelli, T. Reitz, E. Klien, and R. Gregor. TELEIOS: A Database-Powered Virtual Earth Observatory. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Istanbul, Turkey, August 2012. *Demo*. (Same as [d.18].)
- [e.38] G. Graefe, F. Halim, S. Idreos, H. Kuno, and **S. Manegold**. Concurrency Control for Adaptive Indexing. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Istanbul, Turkey, August 2012. (Same as [d.14].)
- [e.37] M. Ivanova, M. L. Kersten, and **S. Manegold**. Data Vaults: a Symbiosis between Database Technology and Scientific File Repositories. In *Proceedings of the International Conference on Scientific and Statistical Database Management (SSDBM)*, pages 485–494, Chania, Crete, Greece, June 2012.
- [e.36] C. Kontoes, I. Keramitsoglou, I. Papoutsis, D. Michail, T. Herekakis, P. Xofis, M. Koubarakis, K. Kyzirakos, M. Karpathiotakis, C. Nikolaou, M. Sioutis, G. Garbis, S. Vassos, **S. Manegold**, M. L. Kersten, H. Pirk, and M. Ivanova. Operational Wildfire Monitoring and Disaster Management Support Using State-of-the-art EO and Information Technologies. In *Proceedings of the International Workshop on Earth Observation and Remote Sensing Applications (EORSA)*, Shanghai, China, June 2012.
- [e.35] H. Pirk, T. Sellam, **S. Manegold**, and M. L. Kersten. X-device Query Processing by Bitwise Distribution. In *Proceedings of the International Workshop on Data Management on New Hardware (DaMoN)*, Scottsdale, AZ, USA, May 2012.
- [e.34] C. Kontoes, I. Papoutsis, D. Michail, T. Herekakis, M. Koubarakis, K. Kyzirakos, M. Karpathiotakis, C. Nikolaou, M. Sioutis, G. Garbis, S. Vassos, I. Keramitsoglou, **S. Manegold**, M. L. Kersten, and H. Pirk. Wildfire Monitoring via the Integration of Remote Sensing with Innovative Information Technologies. In *Proceedings of the EGU General Assembly*, Vienna, Austria, April 2012.

- [e.33] H. Pirk, **S. Manegold**, and M. L. Kersten. Accelerating Foreign-Key Joins using Asymmetric Memory Channels. In *Proceedings of the International Workshop on Accelerating Data Management Systems Using Modern Processor and Storage Architectures (ADMS)*, Seattle, WA, USA, September 2011.
- [e.32] M. L. Kersten, S. Idreos, **S. Manegold**, and E. Liarou. The Researcher’s Guide to the Data Deluge: Querying a Scientific Database in Just a Few Seconds. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Seattle, WA, USA, August 2011. *Challenges & Visions Track Best Paper Award.*(Same as [d.11].)
- [e.31] S. Idreos, **S. Manegold**, H. Kuno, and G. Graefe. Merging What’s Cracked, Cracking What’s Merged: Adaptive Indexing in Main-Memory Column-Stores. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, Seattle, WA, USA, August 2011. (Same as [d.10].)
- [e.30] G. Graefe, S. Ideros, H. Kuno, and **S. Manegold**. Benchmarking Adaptive Indexing. In *Proceedings of the TPC Technology Conference on Performance Evaluation & Benchmarking (TPC TC)*, pages 169–184, Singapore, September 2010.
- [e.29] R. Abdel Kader, M. van Keulen, P. A. Boncz, and **S. Manegold**. Run-time Optimization for Pipelined Systems. In *Alberto Mendelzon International Workshop on Foundations of Data Management (AMW)*, Buenos Aires, Argentina, May 2010.
- [e.28] R. Abdel Kader, P. A. Boncz, **S. Manegold**, and M. van Keulen. ROX: The Robustness of a Run-time XQuery Optimizer Against Correlated Data. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE)*, pages 1185–1188, Long Beach, CA, USA, March 2010. *Demo.*
- [e.27] **S. Manegold**, M. L. Kersten, and P. A. Boncz. Database Architecture Evolution: Mammals Flourished long before Dinosaurs became Extinct. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pages 1648–1653, Lyon, France, August 2009. *10-year Best Paper Award paper for [e.07].* (Same as [d.08].)
- [e.26] R. Abdel Kader, P. A. Boncz, **S. Manegold**, and M. van Keulen. ROX: Run-time Optimization of XQueries. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, pages 615–626, Providence, RI, USA, June 2009.
- [e.25] S. Idreos, M. L. Kersten, and **S. Manegold**. Self-organizing Tuple Reconstruction in Column-stores. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, pages 297–308, Providence, RI, USA, June 2009.
- [e.24] F. Groffen, M. L. Kersten, and **S. Manegold**. Optimising Client Accesses within Armada. In *Workshop on Dependable Distributed Data Management (WDDDM)*, Nuremberg, Germany, March 2009.
- [e.23] L. Sidirourgos, R. Gonçalves, M. L. Kersten, N. Nes, and **S. Manegold**. Column-Store Support for RDF Data Management: not all swans are white. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pages 1553–1563, Auckland, New Zealand, September 2008. (Same as [d.06].)
- [e.22] S. Idreos, M. L. Kersten, and **S. Manegold**. Updating a Cracked Database. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, pages 413–424, Beijing, China, June 2007.
- [e.21] F. Groffen, M. L. Kersten, and **S. Manegold**. Armada: a Reference Model for an Evolving Database System. In *Proceedings of the GI-Fachtagung für Datenbanksysteme in Business, Technologie und Web (BTW)*, pages 417–435, Aachen, Germany, March 2007.
- [e.20] S. Idreos, M. L. Kersten, and **S. Manegold**. Database Cracking. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*, pages 68–78, Asilomar, CA, USA, January 2007.
- [e.19] **S. Manegold**. An Empirical Evaluation of XQuery Processors. In *Proceedings of the International Workshop on Performance and Evaluation of Data Management Systems (ExpDB)*, pages 40–53, Chicago, IL, USA, June 2006. *Updated and extended version available as [d.04] and [f.17].*
- [e.18] P. A. Boncz, T. Grust, M. van Keulen, **S. Manegold**, J. Rittinger, and J. Teubner. MonetDB/XQuery: A Fast XQuery Processor Powered by a Relational Engine. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, pages 479–490, Chicago, IL, USA, June 2006.
- [e.17] P. A. Boncz, J. Flokstra, T. Grust, M. van Keulen, **S. Manegold**, K. S. Mullender, J. Rittinger, and J. Teubner. MonetDB/XQuery—Consistent & Efficient Updates on the Pre/Post Plane. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, number 3896 in Lecture Notes in Computer Science, pages 1190–1193, Munich, Germany, March 2006. *Demo.*

- [e.16] P. A. Boncz, T. Grust, M. van Keulen, **S. Manegold**, J. Rittinger, and J. Teubner. Pathfinder: XQuery—The Relational Way. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pages 1322–1325, Trondheim, Norway, September 2005. Demo.
- [e.15] P. A. Boncz, **S. Manegold**, and J. Rittinger. Updating the Pre/Post Plane in MonetDB/XQuery. In *Proceedings of the International Workshop on XQuery Implementation, Experience and Perspectives (XIME-P)*, Baltimore, MD, USA, June 2005.
- [e.14] M. L. Kersten and **S. Manegold**. Cracking the Database Store. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*, pages 213–224, Asilomar, CA, USA, January 2005.
- [e.13] **S. Manegold**, P. A. Boncz, N. Nes, and M. L. Kersten. Cache-Conscious Radix-Declasser Projections. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pages 684–695, Toronto, ON, Canada, August 2004.
- [e.12] A. R. Schmidt, **S. Manegold**, and M. L. Kersten. Integrated Querying of XML Data in RDBMSs. In *Proceedings of the ACM Symposium on Applied Computing (SAC)*, pages 509–514, Melbourne, FL, USA, March 2003.
- [e.11] **S. Manegold**, P. A. Boncz, and M. L. Kersten. Generic Database Cost Models for Hierarchical Memory Systems. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pages 191–202, Hong Kong, China, August 2002.
- [e.10] M. L. Kersten, **S. Manegold**, P. A. Boncz, and N. Nes. Macro- and Micro- Parallelism in a DBMS. In *Proceedings of the European Conference on Parallel Processing (EuroPar)*, number 2150 in Lecture Notes in Computer Science, pages 6–15, Manchester, UK, August 2001.
- [e.09] **S. Manegold**, P. A. Boncz, and M. L. Kersten. What happens during a Join? — Dissecting CPU and Memory Optimization Effects. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pages 339–350, Cairo, Egypt, September 2000.
- [e.08] **S. Manegold**, A. Pellenkoft, and M. L. Kersten. A Multi-Query Optimizer for Monet. In *Proceedings of the British National Conference on Databases (BNCOD)*, number 1832 in Lecture Notes in Computer Science, pages 36–51, Exeter, United Kingdom, July 2000.
- [e.07] P. A. Boncz, **S. Manegold**, and M. L. Kersten. Database Architecture Optimized for the New Bottleneck: Memory Access. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pages 54–65, Edinburgh, Scotland, UK, September 1999. Received 10-year Best Paper Award at VLDB 2009 [d.08,e.27,h.14].
- [e.06] **S. Manegold**, F. Waas, and M. L. Kersten. Transparent Parallelism in Query Execution. In *Proceedings of the International Conference on Management of Data (COMAD)*, pages 217–234, Hyderabad, India, December 1998.
- [e.05] **S. Manegold** and J. K. Obermaier. Efficient Resource Utilization in Shared-Everything Environments. In *Proceedings of the International Workshop on Issues and Applications of Database Technology (IADT)*, pages 209–216, Berlin, Germany, July 1998.
- [e.04] **S. Manegold** and F. Waas. Thinking Big in a Small World — Efficient Query Execution on Small-scale SMPs. In Jonathan Schaeffer, editor, *Proceedings of the International Symposium on High Performance Computing Systems and Applications (HPCS)*, chapter 14, pages 133–146. Kluwer Academic Publishers, Edmonton, AL, Canada, May 1998.
- [e.03] **S. Manegold**, F. Waas, and D. Gudlat. In Quest of the Bottleneck - Monitoring Parallel Database Systems. In *Proceedings of the European PVM-MPI Users' Group Meeting*, number 1332 in Lecture Notes in Computer Science, pages 277–284, Cracow, Poland, November 1997.
- [e.02] **S. Manegold**, J. K. Obermaier, and F. Waas. Load Balanced Query Evaluation in Shared-Everything Environments. In *Proceedings of the European Conference on Parallel Processing (EuroPar)*, number 1300 in Lecture Notes in Computer Science, pages 1117–1124, Passau, Germany, August 1997.
- [e.01] **S. Manegold**. Efficient Data-Parallel Query Processing in Shared-Everything Environments. In *Technologie für Parallele Datenbanksysteme — Bericht zum Workshop*, München, Germany, February 1997. Technische Universität München, Institut für Informatik. In German.

f. Technical Reports

- [f.17] **S. Manegold**. An Empirical Evaluation of XQuery Processors. Technical Report INS-E0704, CWI, Amsterdam, The Netherlands, March 2006. *Updated and extended version of [e.19]*.
- [f.16] F. Groffen, M. L. Kersten, and **S. Manegold**. Armada: A Reference Model for an Evolving Database System. Technical Report INS-E0603, CWI, Amsterdam, The Netherlands, May 2006.
- [f.15] P. A. Boncz, T. Grust, M. van Keulen, **S. Manegold**, J. Rittinger, and J. Teubner. Loop-lifted Staircase Join: from XPath to XQuery. Technical Report INS-E0510, CWI, Amsterdam, The Netherlands, May 2005.
- [f.14] P. A. Boncz, **S. Manegold**, and J. Rittinger. Updating the Pre/Post Plane in MonetDB/XQuery. Technical Report INS-E0506, CWI, Amsterdam, The Netherlands, April 2005.
- [f.13] P. A. Boncz, T. Grust, **S. Manegold**, J. Rittinger, and J. Teubner. Pathfinder: Relational XQuery Over Multi-Gigabyte XML Inputs In Interactive Time. Technical Report INS-E0503, CWI, Amsterdam, The Netherlands, March 2005.
- [f.12] M. L. Kersten and **S. Manegold**. Cracking the database store. Technical Report INS-E0408, CWI, Amsterdam, The Netherlands, September 2004.
- [f.11] **S. Manegold**, P. A. Boncz, N. Nes, and M. L. Kersten. Cache-Conscious Radix-Decluster Projections. Technical Report INS-E0406, CWI, Amsterdam, The Netherlands, June 2004.
- [f.10] **S. Manegold**, P. A. Boncz, and M. L. Kersten. Generic Database Cost Models for Hierarchical Memory Systems. Technical Report INS-R0203, CWI, Amsterdam, The Netherlands, March 2002.
- [f.09] **S. Manegold**, A. Pellenkoft, and M. L. Kersten. A Multi-Query Optimizer for Monet. Technical Report INS-R0002, CWI, Amsterdam, The Netherlands, January 2000.
- [f.08] **S. Manegold**, P. A. Boncz, and M. L. Kersten. Optimizing Main-Memory Join On Modern Hardware. Technical Report INS-R9912, CWI, Amsterdam, The Netherlands, October 1999.
- [f.07] **S. Manegold**, F. Waas, and M. L. Kersten. On Optimal Pipeline Processing in Parallel Query Optimization. Technical Report INS-R9805, CWI, Amsterdam, The Netherlands, February 1998.
- [f.06] **S. Manegold** and J. K. Obermaier. Efficient Resource Utilization in Shared-Everything Environments. Technical Report INS-R9711, CWI, Amsterdam, The Netherlands, December 1997.
- [f.05] **S. Manegold**, J. K. Obermaier, and F. Waas. Flexible Query Optimization in Parallel Database Systems (Working Paper). Technical Report HUB-IB-80, Humboldt-Universität zu Berlin, Institut für Informatik, Berlin, Germany, February 1997. *In German*.
- [f.04] **S. Manegold**, J. K. Obermaier, and F. Waas. Load Balanced Query Evaluation in Shared-Everything Environments (Extended Version). Technical Report HUB-IB-70, Humboldt-Universität zu Berlin, Institut für Informatik, Berlin, Germany, September 1996.
- [f.03] **S. Manegold**, J. K. Obermaier, F. Waas, and J.-C. Freytag. Load Balanced Query Evaluation in Shared-Everything Environments. Technical Report HUB-IB-61, Humboldt-Universität zu Berlin, Institut für Informatik, Berlin, Germany, June 1996.
- [f.02] **S. Manegold**, J. K. Obermaier, F. Waas, and J.-C. Freytag. Data Threaded Query Evaluation in Shared-Everything Environments. Technical Report HUB-IB-58, Humboldt-Universität zu Berlin, Institut für Informatik, Berlin, Germany, April 1996.
- [f.01] **S. Manegold**, C. Hörner, and A. Heuer. A Look Back on IRIS — Final Report of a Database Research Project. Technical Report Informatik-Bericht 93/7, Technische Universität Clausthal, Institut für Informatik, Clausthal-Zellerfeld, Germany, July 1993. *In German*.

Tutorials, Seminars, Panels, Presentations, Lectures, Courses, in the Media

g. Tutorials, Seminars, & Panels

- [g.06] K. Bereta, M. Koubarakis, **S. Manegold**, G. Stamoulis, and B. Demir. From Big Data to Big Information and Big Knowledge: The Case of Earth Observation Data. In *International Conference on Information and Knowledge Management (CIKM)*, Turin, Italy, October 2018. *Seminar/Tutorial*.

- [g.05] K. Bereta, M. Koubarakis, **S. Manegold**, C. Nikolaou, and G. Stamoulis. A Data Science Pipeline for Linked Earth Observation Data. In *Hellenic Data Management Symposium (HDMS)*, Larnaca, Cyprus, July 2018. *Seminar/Tutorial*.
- [g.04] S. Idreos, **S. Manegold**, and G. Graefe. Adaptive Indexing in Modern Database Kernels. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, pages 566–569, Berlin, Germany, March 2012. *Seminar/Tutorial*.
- [g.03] **S. Manegold** and I. Manolescu. Performance Evaluation in Database Research: Principles and Experience. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, page 1156, St. Petersburg, Russia, March 2009. *Seminar/Tutorial*.
- [g.02] I. Manolescu and **S. Manegold**. Performance Evaluation in Database Research: Principles and Experience. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE)*, Cancun, Mexico, April 2008. *Seminar/Tutorial*.
- [g.01] I. Manolescu and **S. Manegold**. Performance Evaluation and Experimental Assessment — Conscience or Curse of Database Research? In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pages 1441–1442, Vienna, Austria, September 2007. *Panel*.

h. Invited Presentations & Lectures

- [h.41] **S. Manegold**. *Progressive Indices: Adaptive Indexing Without Prejudice*. Humboldt-Universität zu Berlin, Institut für Informatik, Berlin, Germany, April 2018. *Invited Talk*.
- [h.40] **S. Manegold**. *Data Streams and Low Latency Processing*. InfoStrada / Gracernote, Nieuwegein, The Netherlands, December 2016. *Invited Talk*.
- [h.39] **S. Manegold**. *Data Management for Data Science*. InfoStrada / Gracernote, Nieuwegein, The Netherlands, December 2016. *Invited Talk*.
- [h.38] **S. Manegold**. *MonetDB: Open-source Columnar Database Technology Beyond Textbooks*. Universität Rostock, Institut für Informatik, Rostock, Germany, December 2014. *Invited Talk*.
- [h.37] **S. Manegold**. *Performance Evaluation in Database Research: Principles and Experiences*. Universität Rostock, Institut für Informatik, Rostock, Germany, December 2014. *Invited Talk*.
- [h.36] **S. Manegold**. *Data Management for Data Science*. Universität Rostock, Institut für Informatik, Rostock, Germany, December 2014. *Invited Talk*.
- [h.35] **S. Manegold**. *Data Management for Data Science*. De Leidsche Flesch, Studievereniging voor Natuurkunde, Sterrenkunde, Wiskunde & Informatica, Universiteit Leiden, Leiden, The Netherlands, September 2014. *Invited Talk*.
- [h.34] **S. Manegold**. *Data Management for Data Science*. Leiden University, Faculty of Science, Leiden, The Netherlands, September 2014. *Invited Talk*.
- [h.33] **S. Manegold**. *Main-Memory Database Systems*. IEEE International Conference on Data Engineering (ICDE), Chicago, IL, USA, April 2014. *Invited panellist*.
- [h.32] **S. Manegold** and H. Pirk. *Hardware-Conscious Cost Modelling through the Ages*. ETH Zurich, Systems Group, Zurich, Switzerland, September 2013. *Invited Talk*.
- [h.31] **S. Manegold** and H. Mühleisen. *MonetDB: Open-source Columnar Database Technology for Data Analytics*. INRIA Unité de Recherche Saclay - Île-de-France, AVIZ Lab., Orsay Cedex, France, June 2013. *Invited Talk*.
- [h.30] **S. Manegold**. *MonetDB: Open-source Columnar Database Technology Beyond Textbooks*. Universiteit Leiden, Leiden Institute of Advanced Computer Science (LIACS), Leiden, The Netherlands, February 2013. *Invited Talk*.
- [h.29] **S. Manegold**. *Towards Exploratory Data Analysis*. EIT ICT Labs Workshop Big Data, Eindhoven, The Netherlands, December 2012. *Invited Talk*.
- [h.28] **S. Manegold**. *Impact of the convergence of three worlds: High-Performance Computing, Databases, and Analytics*. International Workshop on Accelerating Data Management Systems Using Modern Processor and Storage Architectures (ADMS), Istanbul, Turkey, August 2012. *Invited panellist*.

- [h.27] **S. Manegold.** *Data Intensive Research and Explorative Data Analysis.* Internation Workshop on Big Data Benchmarking (WBDB), San Jose, CA, USA, May 2012. *Invited Talk.*
- [h.26] **S. Manegold.** *Explorative Data Analysis for Data-intensive Research - Not just a Scalability Challenge?* Herbsttreffen der GI-Fachgruppe Datenbanksysteme, Hasso-Plattner-Institut, Potsdam, Germany, November 2011. *Invited Talk. In German.*
[http://www.tele-task.de/archive/lecture/overview/6032/.](http://www.tele-task.de/archive/lecture/overview/6032/)
- [h.25] **S. Manegold.** *In the multi-core age, how do larger, faster, cheaper and more responsive memory sub-systems affect data management?* International Workshop on Accelerating Data Management Systems Using Modern Processor and Storage Architectures (ADMS), Seattle, WA, USA, September 2011. *Invited panellist.*
- [h.24] **S. Manegold.** *Merging What's Cracked, Cracking What's Merged: Adaptive Indexing in Main-Memory Column-Stores.* Technische Universität Dresden, Fakultät für Informatik, Dresden, Germany, September 2011. *Invited Talk.*
- [h.23] **S. Manegold.** *MonetDB: Open-source Columnar Database Technology Beyond Textbooks.* Otto-von-Guericke-Universität Magdeburg, Fakultät für Informatik, Magdeburg, Germany, September 2011. *Invited Talk.*
- [h.22] **S. Manegold.** *The Data Cyclotron Query Processing Scheme.* Technische Universität Ilmenau, Fakultät für Informatik und Automatisierung, Ilmenau, Germany, May 2011. *Invited Talk.*
- [h.21] **S. Manegold.** *MonetDB: Open-source Columnar Database Technology Beyond Textbooks.* Technische Universität Dortmund, Fakultät für Informatik, Dortmund, Germany, February 2011. *Invited Talk.*
- [h.20] **S. Manegold.** *MonetDB: Open-source Columnar Database Technology Beyond Textbooks.* Technische Universität Ilmenau, Fakultät für Informatik und Automatisierung, Ilmenau, Germany, May 2011. *Invited Talk.*
- [h.19] **S. Manegold.** *MonetDB: Open-source Columnar Database Technology Beyond Textbooks.* Technische Universität München, Institut für Informatik, München, Germany, January 2011. *Invited Talk.*
- [h.18] **S. Manegold.** *MonetDB: Open-source Columnar Database Technology in Research and Practice.* HP Labs, Palo Alto, CA, USA, January 2011. *Invited Talk.*
- [h.17] **S. Manegold.** *MonetDB: Open-source Columnar Database Technology in Research and Practice.* IBM Research - Almaden, San Jose, CA, USA, January 2011. *Invited Talk.*
- [h.16] **S. Manegold.** *Performance Evaluation in Database Research: Principles and Experience.* Technische Universität Berlin, Institut für Softwaretechnik und Theoretische Informatik, Berlin, Germany, January 2010. *Invited Talk.*
- [h.15] **S. Manegold.** *MonetDB: Open-source Database Technology Beyond Textbooks.* Humboldt-Universität zu Berlin, Institut für Informatik, Berlin, Germany, January 2010. *Invited Talk.*
- [h.14] **S. Manegold, M. L. Kersten, and P. A. Boncz.** *Database Architecture Evolution: Mammals Flourished long before Dinosaurs became Extinct.* International Conference on Very Large Data Bases (VLDB), Lyon, France, September 2009. *10-year Best Paper Award Keynote [d.08,e.27].*
- [h.13] **S. Manegold.** *High-performance Database Technology beyond Textbooks — Forget (almost) everything you were taught at school/university.* Hochschule Heilbronn, Fakultät für Informatik, Heilbronn, Germany, January 2008. *Invited Talk.*
- [h.12] **S. Manegold.** *MonetDB/XQuery: A High-performance XQuery Processor Powered by a Relational Engine.* Renmin University of China, School of Information, Department of Computer Science and Technology, Beijing, China, June 2007. *Invited Talk.*
- [h.11] **S. Manegold.** *The MonetDB Database Kernel: Architecture & Insights.* Renmin University of China, School of Information, Department of Computer Science and Technology, Beijing, China, June 2007. *Invited Talk.*
- [h.10] P. A. Boncz, T. Grust, M. van Keulen, **S. Manegold**, J. Rittinger, and J. Teubner. *Pathfinder: XQuery—The Relational Way.* IEEE International Conference on Multimedia & Expo (ICME), Amsterdam, The Netherlands, July 2005. *Poster & Demo.*
- [h.09] **S. Manegold.** *Releasing MonetDB/XQuery — Technology Preview 1.* Holland Open Software Conference, Amsterdam, The Netherlands, May 2005. *Invited Talk.*

- [h.08] **S. Manegold.** *High-Performance Database Technology made in Holland.* University of Massachusetts Amherst, Department of Computer Science, Knowledge Discovery Laboratory, Amherst, MA, USA, August 2004. *Invited Talk.*
- [h.07] **S. Manegold.** *High-Performance Database Technology for Hierarchical Memory Systems.* Free University of Bozen-Bolzano, Faculty of Computer Science, Bolzano/Bozen, Italy, June 2004. *Invited Talk.*
- [h.06] **S. Manegold.** *The Impact of Modern Hardware Developments on Database Performance – or – The (Fairy) Tale of “Random Access” Memory.* SAP AG, T-REX Group, Walldorf, Germany, May 2004. *Invited Talk.*
- [h.05] **S. Manegold.** *Database Research at CWI.* SAP AG, T-REX Group, Walldorf, Germany, April 2004. *Invited Talk.*
- [h.04] **S. Manegold.** *The Impact of Modern Hardware Developments on Database Performance – or – The (Fairy) Tale of “Random Access” Memory.* University of Konstanz, Department of Computer & Information Science, Konstanz, Germany, August 2003. *Invited Talk.*
- [h.03] **S. Manegold.** *The Impact of Modern Hardware Developments on Database Performance – or – The (Fairy) Tale of “Random Access” Memory.* Aalborg University, Department of Computer Science, Aalborg, Denmark, July 2003. *Invited Talk.*
- [h.02] **S. Manegold.** *A Generic Cost Model for Parallel Database Systems.* CWI, Amsterdam, The Netherlands, August 1997. *Invited Talk.*
- [h.01] **S. Manegold.** *Nested Transactions in Objectbases.* Humboldt-Universität zu Berlin, Institut für Informatik, Berlin, Germany, October 1995. *Invited Talk.*

i. Courses

- [i.05] **S. Manegold.** *Advanced Data Management for Data Analysis.* Leiden University, Leiden, The Netherlands, autumn 2019–. *Course (Master CS).*
- [i.04] **S. Manegold** and E. Bakker. *Databases & Data Mining.* Leiden University, Leiden, The Netherlands, autumn 2016–2018. *Course (Master CS).*
- [i.03] **S. Manegold.** *Information Systems for Big Data.* Amsterdam Business School, Universiteit van Amsterdam, Amsterdam, The Netherlands, spring 2016–2017. *Course (MBA & MSc “Big Data & Business Analytics”).*
- [i.02] **S. Manegold** and M. L. Kersten. *Advanced Database Techniques.* Universiteit van Amsterdam, Amsterdam, The Netherlands, autumn 2006–2011. *Course (Master CS).*
- [i.01] P. A. Boncz, **S. Manegold**, and M. L. Kersten. *Database Techniques.* Universiteit van Amsterdam, Amsterdam, The Netherlands, spring 2004–2010. *Course (Bachelor CS).*

j. In the Media

- [j.02] **S. Manegold.** *Als je data opslaat op een harde schijf, wordt die dan ook zwaarder? - Hoeveel wegen eentjes en nulletjes? Hoe?Zo!* Radio, June 2010. *Kettingvraag. In Dutch.*
<http://www.cwi.nl/intranet/stefan-manegold-in-hoezo-radio>.
- [j.01] P. A. Boncz, M. L. Kersten, and **S. Manegold.** *Oogstjaar voor CWI-onderzoekers 'We investeren bewust in technologie'.* IPN - ICT-onderzoek Platform Nederland, I/O Magazine, June 2010. *Interview by Daphne Riksen. In Dutch.*
http://www.ictonderzoek.net/3/assets/File/Archief%20Magazine/IO_2010_juni.pdf.