

Curriculum Vitae – Damianos Chatziantoniou

1. Personal Information

Damianos Chatziantoniou.

Address: Pentelis 88, Marousi, Athens 15126, Greece

Phone: +30 (210) 613-7689, *Cell:* +30 (6945) 878-178

Email: damianos@aueb.gr, damianos@cs.columbia.edu

Homepage: <http://www.dmst.aueb.gr/damianos>

2. Education and Awards

- *Department of Computer Science, Columbia University*
 - Ph.D. in Computer Science (June 1997)
 - M.Phil. in Computer Science (February 1996)
 - Major in Database Systems (Data Warehousing, Decision Support Systems, OLAP)
- *Department of Computer Science, Courant Institute of Mathematical Sciences, New York University*
 - M.Sc. in Computer Science (December 1992)
 - Major in Database and Distributed Systems, Minor in Artificial Intelligence
 - GPA : 3.92 (out of 4)
- *Department of Mathematics, University of Athens*
 - B.Sc. in Mathematics (summa cum laude, June 1991)
 - Major in Applied Mathematics

Fellowship for two years of undergraduate studies from the Greek Foundation of Fellowships.

Fellowship from Columbia University for the entire duration of PhD studies.

Fellowship from New York State for part of the M.S. studies at N.Y.U.

3. Academic Experience

- Athens University of Business and Economics
 - *Director of MS in Business Analytics* Sept. 14 – present
 - *Dept. of Management Science and Technology (Associate Prof, tenured)* Dec. 13 – present
 - *Dept. of Management Science and Technology (Assistant Prof)* May 08 – Dec. 13
 - *Dept. of Management Science and Technology (Visiting Prof)* Jan. 02 – May 08
- Technical University of Crete
 - *Electrical and Computing Engineering Dept. (Visiting Prof)* Sept. 02 – Jan. 03
- Stevens Institute of Technology
 - *Computer Science Department (Tenure-Track Assistant Prof)* Sept. 97 – Dec. 99
- Columbia University
 - *Computer Science Department (Adjunct Professor)* Jan. 95 – July 98
 - *Computer Science Department (Teaching Assistant)* Jan. 93 – Dec. 95
- New York University
 - *Computer Science Department (Teaching Assistant)* Sept. 92 – Dec. 92

4. Research Interests

- *Big Data Systems*
 - High-level data modeling, query languages, performance, interoperability, connectivity, integration.
- *Business Intelligence, On-Line Analytical Processing, Data Warehousing.*
 - Modeling, query languages, evaluation algorithms for complex/ad hoc OLAP, systems.
- *Data Streams*
 - Modeling, query languages, evaluation algorithms, systems, relational-streams interoperability.

- *Data Modeling*
Multi-dimensional conceptual models, algebraic forms and languages for complex analytics, object-oriented and attribute-oriented models, schema evolution.
- *Medical Informatics and Clinical Warehouses*
OLAP techniques applied on large medical informatics applications, clinical data warehouses.

5. Publications

Refereed Journals:

1. “*Report on the 8th International Workshop on Business Intelligence for the Real-Time Enterprise*”, Malú Castellanos, Umeshwar Dayal, Nesime Tatbul, Damianos Chatziantoniou, Qiming Chen, **ACM SIGMOD Record, ACM Vol.44(2), p. 62-63, 2015.**
2. “*Tagged MapReduce: Efficiently Computing Multi-analytics Using MapReduce*”, Andreas Williams, Pavlos Mitsoulis-Ntompos, Damianos Chatziantoniou, **International Conference on Data Warehousing and Knowledge Discovery (DaWaK), Lecture Notes in Computer Science, Springer, Vol. 1377, p. 240-251, 2011.**
3. “*Supporting Real-time Supply Chain Decisions Based on RFID Data Streams*”, Damianos Chatziantoniou, Katerina Pramataris, Yannis Sotiropoulos, **Journal of Systems and Software, Elsevier, Volume 84(4), p.700-710, 2011.**
4. “ *θ -Constrained Multi-Dimensional Aggregation*”, Michael Akinde, Michael Bohlen, Damianos Chatziantoniou, Johann Gamper **Information Systems Journal, Elsevier, Volume 36(2), p. 341-358, 2011.**
5. “*ASSET Queries: A Declarative Alternative to MapReduce*”, Damianos Chatziantoniou, Elias Tzortzakakis, **ACM SIGMOD Record, ACM, Vol.38(2), p. 35-42, 2009.**
6. “*ASSET Queries: A Set-Oriented and Column Wise Approach to Modern OLAP*”, Damianos Chatziantoniou, Yannis Sotiropoulos, **International Conference on Very Large Databases (VLDB) 2009, Enabling Business Intelligence for the Real-Time Enterprise Workshop (BIRTE), Lecture Notes on Business Information Processing, Vol.41, p. 66-83, 2010.**
7. “*A Session-Oriented Approach in Modeling Hierarchies of Streams*”, Damianos Chatziantoniou, Achilleas Anagnostopoulos, **Journal of Software: Practice & Experience, Wiley, Vol.38(1), p. 95-115, 2008.**
8. “*Using Grouping Variables to Express Complex Decision Support Queries*”, Damianos Chatziantoniou, **Journal of Data&Knowledge Engineering, Elsevier, Vol.61(1), p. 114-136, 2007.**
9. “*Partitioned Optimization of Complex Queries*”, Damianos Chatziantoniou, Kenneth Ross, **Information Systems Journal, Elsevier, Vol.32(2), p. 248-282, 2007.**
10. “*Extracting and Summarizing Complex Structures in Relational DBMS*”, Damianos Chatziantoniou, **Intern. Journal of Knowledge and Learning, InderScience, Vol.3(4), 2007.**
11. “*Decision Support Queries on a Tape-Resident Data Warehouse*”, Damianos Chatziantoniou, Theodore Johnson, **Information Systems Journal, Elsevier, Vol.30(2), p. 133-149, 2005.**
12. “*NGCE - Network Graphs for Computer Epidemiologists*”
Vasileios Vlachos, Vassiliki Vouzi, Damianos Chatziantoniou, Diomidis Spinellis, **10th Panhellenic Conference on Informatics (PCI), Lecture Notes in Computer Science, Springer, Vol. 3746, p. 672-683, 2005.**
13. “*NESTREAM: Querying Nested Streams*”, Damianos Chatziantoniou, Achilleas Anagnostopoulos, **ACM SIGMOD Record, Vol.33(3), p. 71-78, 2004.**
14. “*SQL Extensions to Enable Decision Support within Database Systems*”, Damianos Chatziantoniou, George Doukidis, **Journal of Decision Systems, Lavoisier, Vol.13(2), p. 155-172, 2004.**
15. “*EMF SQL: Extending SQL to Facilitate Complex OLAP*”, Damianos Chatziantoniou,

- Journal of Business Intelligence, TDWI Research, Vol .8(4), p. 46-55, 2003.**
16. “*Joining Very Large Data Sets*”,
Theodore Johnson, Damianos Chatziantoniou,
**International Conference on Very Large Databases (VLDB),
Databases in Telecommunications Workshop,
Lecture Notes in Computer Science, Springer, Vol. 1819, p. 118-132, 2000.**
 17. “*Extended SQL for Manipulating Clinical Warehouse Data*”,
Stephen Johnson, Damianos Chatziantoniou,
Journal of American Medical Informatics Association (AMIA), Supp(S), p. 819-823, 1999.
 18. “*Complex Aggregation at Multiple Granularities*”,
Kenneth Ross, Divesh Srivastava, Damianos Chatziantoniou,
**International Conference on Extending the Database Technology (EDBT),
Lecture Notes in Computer Science, Springer, Vol. 1377, p. 263-277, 1998.**

Books:

19. “*Incorporating Data Stream Analysis into Decision Support Systems*”
Damianos Chatziantoniou, George Doukidis,
Encyclopedia of Information Science and Technology III, Idea Group Publ., 2005.

Refereed International Conferences:

20. “*Enabling Global Big Data Computations (Vision Paper)*”,
Damianos Chatziantoniou, Panos Louridas,
**International Conference in Extending Database Technology (EDBT) 2018,
Design, Optimization, Languages and Analytical Processing of Big Data (DOLAP) Intl
Workshop.**
21. “*The Data Management Entity: A Simple Abstraction to Facilitate Big Data Systems Interoperability*”,
Damianos Chatziantoniou, Florents Tselai,
**EDBT/ICDT Workshops 2016,
Multi-Engine Data Analytics (MEDAL) Workshop.**
22. “*Introducing Data Connectivity in Big Data Systems*”,
Damianos Chatziantoniou, Florents Tselai,
**International Conference on Management of Data (ACM SIGMOD) 2014,
Data Analytics in the Cloud (DanaC) Workshop.**
23. “*LINKViews: An Integration Framework for Relational and Stream Systems*”,
Yannis Sotiropoulos, Damianos Chatziantoniou,
**International Conference on Very Large Databases (VLDB) 2013,
Enabling Business Intelligence for the Real-Time Enterprise (BIRTE) Workshop.**
- “*Tagged MapReduce: Efficiently Computing Multi-analytics Using MapReduce*”,[\[see 1\]](#)
Andreas Williams, Pavlos Mitsoulis-Ntompos, Damianos Chatziantoniou,
Int. Conf. on Data Warehousing & Knowledge Discovery (DaWaK) 2011.
- “*ASSET Queries: A Set-Oriented and Column Wise Approach to Modern OLAP*”,[\[see 5\]](#)
Damianos Chatziantoniou, Yannis Sotiropoulos,
**International Conference on Very Large Databases (VLDB) 2009,
Enabling Business Intelligence for the Real-Time Enterprise (BIRTE) Workshop.**
24. “*COSTES: Continuous Spreadsheet-Like Computations*”,
Damianos Chatziantoniou, Katerina Pramataris, Yannis Sotiropoulos,
**IEEE International Conference on Data Engineering (ICDE) 2008,
RFID Data Management (RFDm) Workshop.**
25. “*Stream Variables: A Quick but not Dirty SQL Extension for Continuous Queries*”,
Damianos Chatziantoniou, Yannis Sotiropoulos,
**IEEE International Conference on Data Engineering (ICDE) 2007,
Ambient Intelligence, Media and Sensing (AIMS) Workshop.**
26. “*Hierarchical Stream Aggregates: Querying Nested Stream Sessions*”,
Damianos Chatziantoniou, Achilleas Anagnostopoulos,
IEEE Intern. Conf. on Scientific and Statistical DBMS (SSDBM) 2004.
27. “*The MD-Join: An Operator for Complex OLAP*”,

- Damianos Chatziantoniou, Michael Akinde, Theodore Johnson, Samuel Kim,
IEEE International Conference on Data Engineering (ICDE) 2001.
28. “*Extending Complex Ad Hoc OLAP*”,
 Theodore Johnson, Damianos Chatziantoniou,
ACM Conf. on Information and Knowledge Management (CIKM) 1999.
- “*Joining Very Large Data Sets*”, [*see 15*]
 Theodore Johnson, Damianos Chatziantoniou,
**International Conference on Very Large Databases (VLDB)
 Databases in Telecommunications Workshop.**
29. “*The PanQuery Tool and EMF SQL for Complex Data Management*”,
 Damianos Chatziantoniou,
ACM Int. Conf. on Knowledge Discovery and Data Mining (SIGKDD) 1999.
30. “*Evaluation of Ad Hoc OLAP : In-Place Computation*”,
 Damianos Chatziantoniou,
IEEE Intern. Conf. on Scientific and Statistical DBMS (SSDBM) 1999.
31. “*Ad-Hoc OLAP : Expression and Evaluation*”,
 Damianos Chatziantoniou,
IEEE International Conference on Data Engineering (ICDE) 1999.
- “*Complex Aggregation at Multiple Granularities*”, [*see 17*]
 Kenneth Ross, Divesh Srivastava, Damianos Chatziantoniou,
Intern. Conf. on Extending the Database Technology (EDBT) 1998.
32. “*Groupwise Processing of Relational Queries*”,
 Damianos Chatziantoniou, Kenneth Ross,
International Conference in Very Large Databases (VLDB) 1997.
33. “*Querying Multiple Features of Groups in Relational Databases*”,
 Damianos Chatziantoniou, Kenneth Ross,
International Conference in Very Large Databases (VLDB) 1996.

Technical Reports:

34. “*Ad-Hoc OLAP : Expression and Evaluation*”,
 Damianos Chatziantoniou,
Technical Report, Stevens Inst. of Technology, Computer Science Department, 1999.
35. “*Complex Aggregation at Multiple Granularities*”,
 Kenneth Ross, Divesh Srivastava, Damianos Chatziantoniou,
Technical Report, AT&T Research Labs, 1998.
36. “*Querying Multiple Features of Groups in Relational Databases*”,
 Damianos Chatziantoniou, Kenneth Ross,
Technical Report, Columbia University, Computer Science Department, 1996.

6. Research Activities

6.1. Reviewer and Program Committee Member

- Program Committee Member:
 - 34th International Conf. on Very Large Databases (VLDB), 2008.
 - 11th International Conf. in Advances in Database and Information Systems (ADBIS), 2007.
- Reviewer at Refereed International Journals/Conferences:
 - International Conference on Very Large Databases (VLDB) 1998.
 - National Science Foundation (NSF), Information & Intelligent Systems, 2001, 2003.
 - The Very Large Databases Journal (VLDBJ).
 - Information Systems Journal (ISJ).
 - Data & Knowledge Engineering Journal (DKE).
 - ACM Transactions on Database Systems (TODS).
 - IEEE Transactions on Knowledge and Data Engineering (TKDE).
 - Computer Journal, Oxford Press.

6.2. Research Collaborations

- Senior Research Consultant in Query Processing, Aster Data Systems, Redwood City, CA, April 2007 – present. Aster Data Systems was a high-profile technology startup based on Silicon Valley, acquired in 2011 by Teradata.
- AT&T Research Labs, October 1997 – May 1999. Collaboration with Theodore Johnson on OLAP queries over telecom applications.

6.3. International Conference Participation

International Conference on Management of Data (SIGMOD/PODS) 1995-1999, 04, 07, 11, 14.
International Conference on Very Large Databases (VLDB) 1996-2000, 03, 06, 09, 10, 12, 13.
International Conference on Data Engineering (ICDE) 1996, 98, 99, 2001, 07, 08, 09.
International Conference on Extending the Database Technology (EDBT) 1998, 2006, 2014.
International Conference on Knowledge Discovery and Data Mining (KDD) 1999.
Towards an Electronic Patient Record (TEPR) 1995.

6.4. Member of Scientific Communities

ACM, ACM SIGMOD, IEEE.

7. Impact of Research Work

7.1. Selected Citations

There are more than 200 citations on published work (scholar.google.com, Microsoft academic) in major database and data mining textbooks, top-ranked international journals and conferences, and patents (primary references) from IBM, Microsoft, Oracle and NCR. Publication [28] has been included in ACM SIGMOD Digital Review, which presents some of the most influential papers in database literature.

1. Surajit Chaudhuri, Umeshwar Dayal: An Overview of Data Warehousing and OLAP Technology. SIGMOD Record 26(1): 65-74 (1997)
2. Surajit Chaudhuri, Umeshwar Dayal: Data Warehousing and OLAP for Decision Support (Tutorial). SIGMOD Conference 1997: 507-508
3. Surajit Chaudhuri: Review - Querying Multiple Features of Groups in Relational Databases. ACM SIGMOD Digital Review 2: (2000)
4. Divesh Srivastava: Review - Querying Multiple Features of Groups in Relational Databases. ACM SIGMOD Digital Review 2: (2000)
5. Haixun Wang, Carlo Zaniolo: Using SQL to Build New Aggregates and Extenders for Object-Relational Systems. VLDB 2000: 166-175
6. Surajit Chaudhuri, Gerhard Weikum: Rethinking Database System Architecture: Towards a Self-Tuning RISC-Style Database System. VLDB 2000: 1-10
7. Kenneth A. Ross: Review - Providing Better Support for a Class of Decision Support Queries. ACM SIGMOD Digital Review 1: (1999)
8. Jianzhong Li, Doron Rotem, Jaideep Srivastava: Aggregation Algorithms for Very Large Compressed Data Warehouses. VLDB 1999: 651-662
9. Michael O. Akinde, Michael H. Böhlen: Efficient Computation in Subqueries in Complex OLAP. ICDE 2003.
10. H. V. Jagadish, Laks V. S. Lakshmanan, Divesh Srivastava: What can Hierarchies do for Data Warehouses? VLDB 1999: 530-541
11. Shih-Fu Chang, Luis Gravano, Gail E. Kaiser, Kenneth A. Ross, Salvatore J. Stolfo: Database Research at Columbia University. SIGMOD Record 27(3): 75-80(1998)
12. Raghu Ramakrishnan, Donko Donjerkovic, Arvind Ranganathan, Kevin S. Beyer, Muralidhar Krishnaprasad: SRQL: Sorted Relational Query Language. SSDBM 1998: 84-95
13. Luca Cabibbo, Riccardo Torlone: From a Procedural to a Visual Query Language for OLAP. SSDBM 1998: 74-83.

14. Luca Cabibbo, Riccardo Torlone: A Systematic Approach to Multidimensional Databases. SEBD 1997: 361-377.
15. J. Han, S. Chee, and J. Y. Chiang: Issues for On-Line Analytical Mining of Data Warehouses, Proc. of 1998 SIGMOD, Workshop on Research Issues on Data Mining and Knowledge Discovery (DMKD'98), Seattle, Washington, June 1998, pp. 2:1-2:5.
16. Yihong Zhao, Prasad Deshpande, Jeffrey F. Naughton, Amit Shukla: Simultaneous Optimization and Evaluation of Multiple Dimensional Queries. SIGMOD Conference 1998: 271-282
17. Luca Cabibbo, Riccardo Torlone: Querying Multidimensional Databases. DBPL 1997: 319-335
18. Jun Rao, Kenneth A. Ross: Reusing Invariants: A New Strategy for Correlated Queries. SIGMOD Conference 1998: 37-48
19. Michael Jaedicke, Bernhard Mitschang: On Parallel Processing of Aggregate and Scalar Functions in Object-Relational DBMS. SIGMOD Conference 1998: 379-389
20. Weifa Liang, Maria E. Orlowska, Jeffrey X. Yu: Optimizing Multiple Dimensional Queries Simultaneously in Multidimensional Databases. VLDB Journal 8(3-4): 319-338(2000)
21. Prasad Deshpande, Jeffrey F. Naughton: Aggregate Aware Caching for Multi-Dimensional Queries. EDBT 2000: 167-182
22. H. V. Jagadish, Laks V. S. Lakshmanan, Tova Milo, Divesh Srivastava, Dimitra Vista: Querying Network Directories. SIGMOD Conference 1999: 133-144
23. Frédéric Gingras, Laks V. S. Lakshmanan: nD-SQL: A Multi-Dimensional Language for Interoperability and OLAP. VLDB 1998: 134-145
24. Chuck Cranor, Theodore Johnson, Oliver Spatscheck: How To Query Network Traffic Data Using Data Streams. VLDB 2002.
25. Raghu Ramakrishnan, Johannes Gehrke: Database Management Systems (1st, 2nd, 3rd editions). McGraw-Hill.
26. Eike Schallehn, Kai-Uwe Sattler, Gunter Saake : Advanced Grouping and Aggregation for Data Integration. CIKM 2001.
27. Matteo Golfarelli, Dario Maio, Stefano Rizzi: The Dimensional Fact Model: A Conceptual Model for Data Warehouses. Int. J. Cooperative Inf. Syst. 7(2-3): 215-247 (1998).
28. Eike Schallehn, Kai-Uwe Sattler, Gunter Saake: Extensible Grouping and Aggregation for Data Reconciliation. Proc. 4th Int. Workshop on Engineering Federated Information Systems, EFIS'01.
29. Eike Schallehn, Kai-Uwe Sattler, Gunter Saake: Extensible and Similarity-based Grouping for Data Integration . ICDE 2002.
30. Michael O. Akinde, Michael H. Böhlen: Generalized MD-Joins: Evaluation and Reduction to SQL. Databases in Telecommunications 2001.
31. Haixun Wang, Carlo Zaniolo: Database System Extensions for Decision Support: the AXL Approach. ACM SIGMOD 2000 Workshop on Research Issues in Data Mining and Knowledge Discovery
32. Haixun Wang, Carlo Zaniolo: User-Defined Aggregates for Data Mining. ACM SIGMOD 1999 Workshop on Research Issues in Data Mining and Knowledge Discovery
33. Haixun Wang, Carlo Zaniolo: User Defined Aggregates in Object-Relational Systems. ICDE 2000.
34. Datta, Moon: Have your Data and Index it, too. Efficient Storage and Indexing for Data Warehouses 1998.
35. Michael O. Akinde, Michael H. Böhlen, Theodore Johnson, Laks V. S. Lakshmanan, Divesh Srivastava: Efficient OLAP Query Processing in Distributed Data Warehouses. EDBT 2002
36. César A. Galindo-Legaria, Milind Joshi: Orthogonal Optimization of Subqueries and Aggregation. SIGMOD Conference 2001.
37. Matteo Golfarelli, Stefano Rizzi: View materialization for nested GPSJ queries. Second Intl. Workshop on Design and Management of Data Warehouses, DMDW 2000.
38. Michael O. Akinde, Michael H. Böhlen, Theodore Johnson, Laks V. S. Lakshmanan, Divesh Srivastava: Efficient OLAP query processing in distributed data warehouses. Information Systems 28(1-2) (2003)
39. Panos Vassiliadis: Gulliver in the land of data warehousing: practical experiences and observations of a researcher. DMDW 2000.
40. Srikanth Bellamkonda, Tolga Bozkaya, Bhaskar Ghosh, Abhinav Gupta, John Haydu, Sankar Subramanian, Andrew Witkowski: Analytic Functions in Oracle 8i. Stanford presentation/ORACLE technical report. Available at: <http://www-db.stanford.edu/dbseminar/Archive/SpringY2000/speakers/agupta/paper.pdf>.

41. Oracle Corporation: Analytic Functions for Oracle 8i, an Oracle technical white paper, October 1999, otn.oracle.com/products/oracle8i/pdf/8ifunctions.pdf.
42. H. V. Jagadish, Laks V. S. Lakshmanan, Divesh Srivastava, Keith Thompson: TAX: A Tree Algebra for XML. DBPL 2001: 149-164
43. Johannes Gehrke, Flip Korn, Divesh Srivastava: On Computing Correlated Aggregates Over Continual Data Streams. SIGMOD Conference 2001
44. SQL Standards Committee - American National Standards Institute. Information technology – Database languages - SQL - AMENDMENT 1: On-Line Analytical Processing (SQL/OLAP), 2001. Supplement to ISO/IEC 9075:1999.
45. Antonio Badia, Dev Anand: Fighting Redundancy in SQL. DaWaK 2003: 401-411.
46. Ramzi A. Haraty, Roula C. Fany: Query Acceleration in Distributed Database Systems. Revista Colombiana de Computación 2(1): 19-34 (2001)
47. Torsten Grust: Monad Comprehensions. A Versatile Representation for Queries, in P.M.D. Gray, L. Kerschberg, P.J.H. King, A. Poulouvasilis (eds.), The Functional Approach to Data Management - Modeling, Analyzing and Integrating Heterogeneous Data, Springer Verlag, ISBN 3-540-00375-4, September 2003.
48. Sally I. McClean, Bryan W. Scotney, Mary Shapcott: Using Background Knowledge in the Aggregation of Imprecise Evidence in Databases. Data Knowl. Eng. 32(2): 131-143 (2000).
49. Torsten Grust: Comprehending Queries. PhD Thesis. University of Konstanz, Department of Computer and Information Science, Database Systems Research Group.
50. Torsten Grust: Comprehending Queries. In *Ausgezeichnete Informatik-dissertationen 1999* (Distinguished Dissertations in Computer Science), pages 74-83, B.G.Teubner, Sep 2000.
51. Marc Laporte, Noel Novelli, Rosine Cicchetti, Lotfi Lakhal: Computing Full and Iceberg Datacubes Using Partitions. International Symposium on Methodologies for Intelligent Systems (ISMIS) 2002: 244-254.
52. T. Y. Lin, Xiaohua Tony Hu, Eric Louie: Using Attribute Value Lattice to Find Closed Frequent Itemsets. International Conference on Data Mining and Knowledge Discovery Tools, Methods and Techniques (2003).
53. Nick Koudas, Divesh Srivastava: Data Stream Query Processing: A Tutorial. VLDB 2003: 1149
54. Xiaohua Tony Hu, T. Y. Lin, Eric Louie: Bitmap Techniques for Optimizing Decision Support Queries and Association Rule Algorithms. IEEE Seventh International Database Engineering and Applications Symposium (IDEAS'03) (2003).
55. Akinde Michael: Complex and Distributed On-line Analytical Processing. Department of Computer Science, Aalborg University. PhD Thesis.
56. Abraham Silberschatz, Frank Korth and S. Sudarshan: Database System Concepts, 4th edition, McGraw-Hill, ISBN 0-07-228363-7, 2003.
57. J. Abello, P. Pardallos, M. Resende: Handbook of massive data sets, Kluwer Academic Publishers, ISBN 1-4020-0489-3, 2002.
58. Jiawei Han and Micheline Kamber: Data Mining: Concepts and Techniques, Morgan Kaufmann Publishers, ISBN 1-55860-489-8, 2000.
59. Harvey Miller, Jiawei Han: Geographic Data Mining and Knowledge Discovery, Taylor & Francis Pubs, ISBN 0-4152-3369-0, 2001.
60. Tamraparni Dasu, Theodore Johnson: Exploratory Data Mining and Data Cleaning, John Wiley, ISBN 0-4712-6851-8, 2003.
61. Jiawei Han: Handbook of data mining and knowledge discovery, Oxford University Press, ISBN 0-19-511831-6, 2002.
62. Matthias Jarke, Maurizio Lenzerini, Yannis Vassiliou, Panos Vassiliadis: Fundamentals of Data Warehouses, Springer, ISBN 978-3540420897, 2003.
63. Maurizio Rafanelli: Multidimensional Databases: Problems and Solutions, Idea Group 2003, ISBN 1-59140-053-8.
64. Venkatesh Ganti, Johannes Gehrke, Raghu Ramakrishnan: Mining Data Streams under Block Evolution. SIGKDD Explorations 3(2): 1-10 (2002)
65. Norman May, Sven Helmer, Guido Moerkotte: Three Cases for Query Decorrelation in XQuery. Xsym 2003 (1st International XML Database Symposium): 70-84.
66. Antonio Badia, Matthew Niehues: Optimization of Sequences of Relational Queries in Decision-Support Environments. DaWaK 1999: 126-131
67. Ling Feng, Tharam S. Dillon: Using Fuzzy Linguistic Representations to Provide Explanatory Semantics for Data Warehouses. IEEE Trans. Knowl. Data Eng. 15(1): 86-102 (2003).

68. Helen Thomas, Anindya Datta: A Conceptual Model and Algebra for OLAP in Decision Support Databases. *Journal of Information Systems Research (INFORMS)*, March 2001, Vol 12(1): 83-102
69. Betsy L. Humphreys: Electronic Health Record Meets Digital Library, *Journal of American Medical Informatics Association*, September 2000, 7 (5): 444-452.
70. Anhøj Jacob: Generic Design of Web-Based Clinical Databases, *Journal of Medical Internet Research* 2003, 5(4):e27.
71. G. Hripcsac, S. Bakken, P. Stetson, V. Patel. Mining complex clinical data for patient safety research: a framework for event discovery. *Journal of Biomedical Informatics*, 36(1-2):150-160, (2003).
72. Calisto Zuzarte, Hamid Pirahesh, Wenbin Ma, Qi Cheng, Linqi Liu, Kwai Wong: WinMagic : Subquery Elimination Using Window Aggregation. *SIGMOD Conference 2003*: 652-656
73. Charles D. Cranor, Theodore Johnson, Oliver Spatscheck, Vladislav Shkapenyuk: Gigascope: A Stream Database for Network Applications. *SIGMOD Conference 2003*: 647-651
74. Andrew Witkowski, Srikanth Bellamkonda, Tolga Bozkaya, Gregory Dorman, Nathan Folkert, Abhinav Gupta, Lei Sheng, Sankar Subramanian: Spreadsheets in RDBMS for OLAP. *SIGMOD Conference 2003*: 52-63
75. Michael Akinde, Michael Bohlen: Efficient Computation of Subqueries in Complex OLAP. *ICDE Conference 2003*.
76. Surajit Chaudhuri, Raghav Kaushik, Jeffrey F. Naughton: On Relational Support for XML Publishing: Beyond Sorting and Tagging. *SIGMOD Conference 2003*: 611-622
77. Rohit Ananthakrishna, Abhinandan Das, Johannes Gehrke, Flip Korn, S. Muthukrishnan, Divesh Srivastava: Efficient Approximation of Correlated Sums on Data Streams. *TKDE* 15(3): 569-572 (2003)
78. Surajit Chaudhuri, Prasanna Ganesan, Vivek R. Narasayya: Primitives for Workload Summarization and Implications for SQL. *VLDB 2003*: 730-741
79. Norman May, Sven Helmer, Guido Moerkotte: Nested Queries and Quantifiers in an Ordered Context, *ICDE 2004*.
80. Stratos Papadomanolakis, Anastassia Ailamaki: Workload-Driven Schema Design for Large Scientific Databases. *IEEE Data Eng. Bull.* 27(4): 21-28 (2004).
81. Antonio Badia: Computing SQL Queries with Boolean Aggregates. *DaWaK 2003*: 391-400.
82. Stratos Papadomanolakis, Anastassia Ailamaki: AutoPart: Automating Schema Design for Large Scientific Databases Using Data Partitioning. *SSDBM 2004*: 383-392.
83. Kenneth A. Ross, Julia Stoyanovich: Symmetric Relations and Cardinality-Bounded Multisets in Database Systems. *VLDB 2004*: 912-923.
84. Carlos Ordonez: Horizontal aggregations for building tabular data sets. *DMKD 2004*: 35-42.
85. Hyuk-Min Lee, Kyung-Chang Kim: A Strategy for Maintaining Client-Based Web Warehouse. *Human.Society@Internet 2005*: 144-154.
86. Vinay Kanitkar, Alex Delis: Efficient Processing of Client Transactions in Real-Time. *Distributed and Parallel Databases* 17(1): 39-74 (2005).
87. Mohammed I. Rafiq, Martin J. O'Connor, and Amar K. Das: Computational Method for Temporal Pattern Discovery in Biomedical Genomic Databases. *IEEE Computational Systems Bioinformatics Conference (CSB'05)*.
88. Philip J. B. Brown, Vic J. Rayward-Smith: Theoretical Potential of Semantic Based Knowledge Discovery From Terminology Populated Clinical Databases. *Annual Symposium of American Medical Informatics Association (AMIA) 1999*.
89. Nathan Folkert, Abhinav Gupta, Andrew Witkowski, Sankar Subramanian, Srikanth Bellamkonda, Shrikanth Shankar, Tolga Bozkaya, Lei Sheng: Optimizing Refresh of a Set of Materialized Views. *VLDB 2005*: 1043-1054.
90. Michael Bohlen, Johann Gamper, Christian Jensen: Multi-Dimensional Aggregation for Temporal Data. *Advances in Database Technology (EDBT) 2006*.
91. Netz, Amir (Microsoft Corporation). Aggregations performance estimation in database systems. *United States Patent 6374234*.
92. Netz, Amir; Pasumansky, Mosha (Microsoft Corporation). Usage based aggregation optimization. *United States Patent 6438537*.
93. Galindo-Legaria, Cesar, Graefe Goetz, Joshi Milind, Bunker Ross (Microsoft Corp.): System and Method for Segmented Evaluation of Database Queries. *US Patent 6748392*.
94. Berger, Alexander; Netz, Amir; Pasumansky, Mosha (Microsoft Corporation). Record for a multidimensional database with flexible paths. *United States Patent 6446059*.

95. Leung, Ting Yu; Wang, Haixun (IBM Corporation). Query transformation and simplification for group by queries with rollup/grouping sets in relational database management systems. United States Patent 6574623.
96. Kelkar, Bhooshan (IBM Corporation). Systems, methods, and computer program products to rank and explain dimensions associated with exceptions in multidimensional data. United States Patent 6714940.
97. Datta, Anindya (Chutney Technologies, Inc.): Dynamic Page Generation Acceleration Using Component-level Caching. United States Patent 6622168.
98. Deshpande, Prasad Manikarao; Ramasamy, Karthikeyan; Shukla, Amit; Naughton, Jeffrey F. (NCR Corporation). Active caching for multi-dimensional data sets in relational database management system. United States Patent 6601062.
99. Deshpande, Prasad Manikarao; Ramasamy, Karthikeyan; Shukla, Amit; Naughton, Jeffrey F. (NCR Corporation). Method for determining the computability of data for an active multi-dimensional cache in a relational database management system. United States Patent 6763357.
100. Zuzarte, Calisto P.; Cheng, Qi; Ma, Wenbin (IBM Corporation). Method and system for aggregation subquery join elimination. United States Patent 20040220911.
101. Witkowski, Andrew; Bellamkonda, Srikanth; Bozkaya, Tolga; Gupta, Abhinav; Folkert, Nathan; Subramanian, Sankar (Oracle Corporation). Run-time optimizations of queries with SQL spreadsheet. United States Patent 20040133567.
102. Witkowski, Andrew; Bellamkonda, Srikanth; Bozkaya, Tolga; Gupta, Abhinav; Folkert, Nathan; Subramanian, Sankar (Oracle Corporation). Compile-time optimizations of queries with SQL spreadsheet. United States Patent 20070055661.
103. Gudbjartsson, Hakon; Arnarson, Thorvaldur S.; Rovensky, Pavol; Palmason, Vilmundur (deCODE genetics ehf). Set definition language for relational data. United States Patent 20050050030.
104. Lei Chen, Raghu Ramakrishnan, Paul Barford, Bee-Chung Chen, Vinod Yegneswaran: Composite Subset Measures. VLDB 2006: 403-414.
105. Surajit Chaudhuri, Venkatesh Ganti, Raghav Kaushik: A Primitive Operator for Similarity Joins in Data Cleaning. ICDE 2006.
106. Elaheh Pourabbas, Arie Shoshani: The Composite OLAP-Object Data Model: Removing an Unnecessary Barrier. SSDBM 2006: 291-300.
107. Radhika Sridhar, Padmashree Ravindra, Kemafor Anyanwu: RAPID: Enabling Scalable Ad-Hoc Analytics on the Semantic Web. International Semantic Web Conference 2009: 715-730.
108. Graham Cormode, Srikanta Tirthapura, Bojian Xu: Time-Decayed Correlated Aggregates over Data Streams. SIAM International Conference on Data Mining 2009: 269-280.
109. Yufei Tao, Xiaokui Xiao: Efficient temporal counting with bounded error. VLDB J. 17(5): 1271-1292 (2008).
110. Lei Chen, Christopher Olston, Raghu Ramakrishnan: Parallel Evaluation of Composite Aggregate Queries. ICDE 2008: 218-227.
111. Yann Gripay, Frédérique Laforest, Jean-Marc Petit: Towards service-oriented continuous queries in pervasive systems. Ingénierie des Systèmes d'Information 13(5): 33-57 (2008).
112. Elizabeth Van Ruitenbeek, Tod Courtney, William H. Sanders, Fabrice Stevens: Quantifying the Effectiveness of Mobile Phone Virus Response Mechanisms. DSN 2007: 790-800.
113. George Lekakos: Exploiting RFID digital information in enterprise collaboration. Industrial Management and Data Systems 107(8): 1110-1122 (2007).
114. Stephanos Androutsellis-Theotokis, Diomidis Spinellis, Vasileios Vlachos: The MoR-Trust Distributed Trust Management System: Design and Simulation Results. Electr. Notes Theor. Comput. Sci. 179: 3-15 (2007).
115. Antonio Badia, Stijn Vansummeren: Non-linear prefixes in query languages. PODS 2007: 185-194.
116. Yann Gripay, Frédérique Laforest, Jean-Marc Petit: Towards Action-Oriented Continuous Queries in Pervasive Systems. BDA 2007.
117. Minsoo Lee, Sookyung Song, Yun-mi Kim, Hyoseop Shin: Supporting Efficient Grouping and Summary Information for Semistructured Digital Libraries. ICADL 2006: 400-409.
118. Xuepeng Yin, Torben B Pedersen: What Can Hierarchies Do for Data Streams? BIRTE 2006: 4-19.
119. Norman May, Guido Moerkotte: Main Memory Implementations for Binary Grouping. XSym 2005: 162-176.
120. Weili Wu, Hong Gao, Jianzhong Li: New Algorithm for Computing Cube on Very Large Compressed Data Sets. IEEE Trans. Knowl. Data Eng. 18(12): 1667-1680 (2006).

121. Bin Cao, Antonio Badia: Exploiting maximal redundancy to optimize SQL queries. *Knowl. Inf. Syst.* 20(2): 187-220 (2009).
122. Surajit Chaudhuri, Umeshwar Dayal, Venkatesh Ganti: Data management technology for decision support systems. *Advances in Computers* 62: 294-326 (2004).
123. Jian Zhu, Hanshi Wang: Application of e-commerce personality searching based on RSS. *IEEE Information Management and Engineering (ICIME)*, 197-199 (2010).
124. Hyunho Lee, Wonsuk Lee: Query Optimization for Web BBS by Analytic Function and Function-Based Index in Oracle DBMS. *ALPIT 2007*: 606-611.
125. Jussi Rasinmäki: XQuery as a retrieval mechanism for longitudinal multiscale forest resource data. *Environmental Modelling and Software* 24(10): 1153-1162 (2009)
126. Bee-Chung Chen, Raghu Ramakrishnan, Jude W. Shavlik, Pradeep Tamma: Bellwether analysis: Searching for cost-effective query-defined predictors in large databases. *TKDD* 3(1): (2009).
127. David Auber, Noel Novelli, Guy Melançon: Visually Mining the Datacube using a Pixel-Oriented Technique. *11th International Conference Information Visualization (IV '07) 2007*: 3-10.
128. Shichao Zhang, Rifeng Wang, Zhi Jin: Strategies for complex data cube queries. *Appl. Intell.* 31(3): 332-346 (2009).
129. Shichao Zhang, Rifeng Wang, Yanping Guo: Efficient Computation of Multi-feature Data Cubes. *KSEM 2006*: 612-624.
130. Norman May, Sven Helmer, Guido Moerkotte: Strategies for query unnesting in XML databases. *ACM Trans. Database Syst.* 31(3): 968-1013 (2006).
131. Guido Moerkotte, Thomas Neumann: Dynamic programming strikes back. *SIGMOD Conference 2008*: 539-552.
132. Bin Cao, Antonio Badia: SQL query optimization through nested relational algebra. *ACM Trans. Database Syst.* 32(3): 18 (2007).
133. Luping Ding, Elke A. Rundensteiner: Index tuning for parameterized streaming groupby queries. *SSPS 2008*: 68-78.

7.2. Impact on international standards, commercial DBMS and US Patents

- *Microsoft - SQL Server*: Microsoft has implemented part of the research work in “Groupwise Processing of Relational Queries” [27] in the Query Processor of Microsoft SQL Server and issued patents citing this article as one of the primary references [91,92,93,94].
- *Oracle*: The research work in “Querying Multiple Features of Groups in Relational Databases” [28] has greatly influenced ORACLE’s Analytic Functions, currently a benchmark in the Data Warehousing industry. US Patents awarded to Oracle [101,102] cite the above-mentioned article as the primary reference.
- *IBM - DB2*: Part of the research work in “Groupwise Processing of Relational Queries” [27] implemented as an optimization technique in the query processor of DB2. IBM US Patents [95,96] reference this article.
- *Standard SQL*: Research work in “Querying Multiple Features of Groups in Relational Databases” influenced the syntax of SQL (GROUPING SETS and WINDOW constructs) [44].

8. Teaching Experience

Graduate Courses & Advanced Seminars:

- Big Data & Business Analytics (Athens University of Economics and Business, *Advanced Seminar*)
- Business Intelligence (Athens University of Economics and Business, *Advanced Seminar*)
- e-Business Technologies and Systems (Athens University of Economics and Business)
- Advanced Data Structures and Algorithms (Stevens Institute of Technology)
- Databases II (Stevens Institute of Technology)
- Digital Logic (Columbia University)
- Advanced Topics in Database Systems (Columbia University)

Undergraduate Courses:

- Database Information Systems (Athens University of Economics and Business)

- Programming Languages I – Java (Athens University of Economics and Business)
- Data Warehousing and Data Mining (Athens University of Economics and Business)
- Systems Analysis (Athens University of Economics and Business, co-Instructor)
- Operating Systems (Technical University of Crete)
- Algorithms and Complexity (Technical University of Crete)
- Databases I (Stevens Institute of Technology)
- Introduction to Programming (Columbia University)
- Data Structures and Algorithms (Columbia University)

9. Non-Academic Experience

- *Aster Data Inc., Redwood City, CA (Senior Research Consultant)* July 07 – March 08
Aster Data Systems was a high-profile technology startup based in Silicon Valley. Its product is a highly distributed, parallel, large-scale data warehousing system. On April of 2011, Aster Data was acquired by Teradata for \$260M.
- *Panakea Software Inc. (CTO and Founder)* Oct. 98 – Jul. 02
Develop and market patent-pending database querying/reporting technology to make advanced data analysis (decision support/data mining) applications easier and faster. Technology is offered as a product and/or as consulting services. These techniques have been implemented with great success in AT&T, Dun&Bradstreet, etc.
- *VoiceWeb S.A. (Chairman of the Board, CTO and Founder)* Jan. 01 – Jan. 03
One of the largest speech-recognition companies in the south east Europe. Implemented the first voice portal in Greece – and the second in whole Europe. Customers include telecoms, banks, entertainment complexes, pharmaceuticals, etc. It has offices in eight countries.
- *Sybase (Intern)* Summer 94
Client applications research & development
- *Newgate Associates (Consultant)* Jan. 93 – Oct. 93
Financial applications programmer (integrating sources, data analysis, consistency, reconciliation)

10. Systems and Languages

- *Languages:*
Excellent knowledge of: C/C++, Java, SQL, Cobol, Pascal, Ada, Fortran, Lisp, Prolog.
- *Database Systems/Platforms:*
SQL Server, Oracle, Hadoop.
- *Telecom Systems/Platforms:*
Dialogic, Philips/Nuance Speech Engines, Envovx CTI.

11. References

Available upon request