
Curriculum Vitae

Neoklis Polyzotis

Associate Professor

Department of Computer Science

UC Santa Cruz

Phone: (831) 459 1304 Fax: (831) 459 4829 E-mail: alkis@cs.ucsc.edu

URL: <http://www.cs.ucsc.edu/~alkis>

EMPLOYMENT

2009—present	Associate Professor, UC Santa Cruz
June 2010	Visiting Faculty, EPFL, Switzerland
2003—2009	Assistant Professor, UC Santa Cruz
July 2004, July 2005, September 2006, September 2007, July 2008, July 2009	Visiting Scientist, INRIA Futurs, France
1998—2003	Research Assistant, University of Wisconsin at Madison
Summer 2001	Summer Intern, Bell Labs, Lucent Inc.
1996—1997	System and Network Administrator, National Technical Univ. of Athens, Greece

EDUCATION

2003	Ph.D. in Computer Sciences, University of Wisconsin at Madison
2000	M.Sc. in Computer Sciences, University of Wisconsin at Madison
1997	Diploma in Electrical and Computer Engineering, National Technical Univ. of Athens, Greece

AWARDS

2008	Best Newcomer Award for the paper “Evaluating rank joins with optimal cost” in Proceedings of ACM Symposium on Principles of Database Systems (PODS), pp. 43-52, 2008 co-authored with K. Schnaitter
2007	Runner-up for Best Paper Award for the paper “Depth Estimation for Ranking Query Optimization” in Proceedings of International Conf. on Very Large Data Bases (VLDB), pp. 902-913, 2007 co-authored with K. Schnaitter and J. Spiegel

2005 National Science Foundation (NSF) CAREER Award

EXTRAMURAL FUNDING

2010 Department of Energy Award
“Adding Data Management Services to Parallel File Systems”
PI: Scott Brandt, co-PIs: Carlos Maltzahn, Neoklis Polyzotis
Amount: \$1,012,535

2010 National Science Foundation (NSF) Award IIS-1018914
“Novel Techniques for Automated Index Tuning”
PI: Neoklis Polyzotis
Amount: \$499,676

2008 SAP Inc., Research Grant
PI: Scott Brandt, co-PIs: Carlos Maltzahn, Neoklis Polyzotis
Amount: \$100,000

2007-2011 UCSC/Los Alamos Institute for Scalable Scientific Data Management
Amount: \$150,000

2006 IBM Faculty Development Award
Amount: \$30,000

2005 IBM Faculty Development Award
Amount: \$30,000

2005 National Science Foundation (NSF) CAREER Award. IIS-0447966
“Novel Summarization Techniques for Semi-Structured Data”
PI: Neoklis Polyzotis
Amount: \$500,000

2003 Microsoft Corporation, Industrial Gift
Amount: \$25,000

INVITED TALKS

2010

- “Semi-Automatic Index Tuning for Database Systems”, Summer Research Institute, EPFL, Switzerland.

2009

- “Autonomic Index Tuning in Database Systems”, Los Alamos National Laboratories.
- “Efficient and Robust Processing of top-k Relational Queries”, Microsoft Search Labs.
- “Autonomic Index Tuning in Database Systems”, HP Labs.

2008

- “The data ring: Community Content Sharing”, Summer Research Institute, EPFL.
- “Depth Estimation for Ranking Query Optimization”, Database Seminar, University of Waterloo.

- “Depth Estimation for Ranking Query Optimization”, Database Seminar, University of Toronto.

2007

- “Data Statistics for Query Optimization: An Old New Problem”, EPFL.
- “Depth Estimation for Ranking Query Optimization”, IBM Almaden Research Center.

2006

- “Data Statistics for Query Optimization: An Old New Problem”, Database Seminar, UC Berkeley.
- “Graph-Based Synopses for Relational Selectivity Estimation”, Database Seminar, Stanford University.
- “Graph-Based Synopses for Relational Selectivity Estimation”, IBM Almaden Research Center.
- “Graph-Based Synopses for Relational Selectivity Estimation”, Oracle Corp.

2005

- “Approximate Query Answering Techniques for XML Databases”, Gemo Database Seminar, INRIA Futurs, 2005.
- “Approximate Query Answering for XML Databases”, SIGMOD Program Committee Workshop, 2005.
- “XCluster Synopses for Structured XML Content”, IBM Almaden Research Center, 2005.

2004

- “Approximate Query Answering Techniques for XML Databases”, Database Seminar, UC Berkeley, 2004.
- “Approximate Query Answering Techniques for XML Databases”, IBM Almaden Research Center, 2004.
- “Approximate Query Answering Techniques for XML Databases”, Database Seminar, Carnegie Mellon University, 2004.
- “Summarization Techniques for XML Databases”, Sybase Corp., 2004.
- “Approximate XML Query Answers”, Computer Science Colloquium, San Jose State University, 2004.

SERVICE TO THE PROFESSION

- **Program Committee Co-Chair:**
 - IEEE Workshop on Self-Managing Database Systems (SMDB): 2011.
 - International Conference on Very Large Databases (VLDB)–Demonstrations Track: 2009.
 - International Conference on Very Large Databases PhD Workshop: 2008.
 - ACM International Workshop on Web Information and Data Management (WIDM): 2007, 2008.

- **Conference Program Committee Member:**
 - Conference on Innovative Data Systems Research (CIDR): 2011.
 - International Conference on Extending Database Technology (EDBT): 2011.
 - WWW: 2010.
 - ACM SIGMOD International Conference on Management of Data: 2005, 2009, 2011.
 - International Conference on Very Large Databases (VLDB): 2006, 2007, 2010.
 - IEEE International Conference on Data Engineering (ICDE): 2005, 2006, 2008-2011.
 - International Conference on Extending Database Technology (EDBT): 2006, 2010.
 - International Conference on Database Theory (ICDT): 2010.
 - IEEE International Conference on Data Mining (ICDM): 2004, 2005, 2008.
 - ACM SIGKDD International Conference on Knowledge Discovery and Data Mining: 2007.
 - ACM International Conference on Information and Knowledge Management (CIKM): 2006, 2008.
 - International Conference on Management of Data (COMAD): 2005.
 - International Conference on Very Large Databases (VLDB) – Demonstrations Track: 2008.
 - International Conference on Extending Database Technology (EDBT) – Demonstrations Track: 2009.
- **Workshop Program Committee Member:**
 - International Workshop on Data Management on New Hardware (DAMON): 2011.
 - ACM SIGMOD PhD Workshop on Innovative Database Research (IDAR): 2009.
 - International Workshop on Self-Managing Database Systems (SMDB): 2009.
 - International Workshop on Ranking in Databases (DBRANK): 2007, 2008, 2009, 2010.
 - ACM International Workshop on Web Information and Data Management (WIDM): 2006.
 - International XML Database Symposium (XSym): 2004, 2006.
 - International Workshop on XQuery Implementation, Experience and Perspectives (XIME-P), 2005.
- **Proceedings Chair:**
 - ACM SIGMOD International Conference on Management of Data, 2006.
- **External Reviewer:**
 - Conferences: SIGMOD 2004, PODS 2004, SIAMDM 2004, PKDD 2004, CIKM 2004, VLDB 2003.
 - Journals: ACM TODS, IEEE TKDE, ACM TOIT, IS, VLDB Journal, IBM Systems Journal, JSS.

PUBLICATIONS

Chapters in Books

1. “XML Selectivity Estimation”.
M. Ramanath, J. Freire, and N. Polyzotis.
To appear in Encyclopedia of Database Systems (Ling Liu and Tamer Ozsu, Eds.),
Springer.
2. “Very Large Databases”.
M. Garofalakis and N. Polyzotis.
In Wiley Encyclopedia of Computer Science and Engineering (Benjamin W. Wah, Ed.),
John Wiley & Sons, Inc., 2008.

Articles in Scholarly Journals

3. “Optimal Algorithms for Evaluating Rank Joins in Database Systems”.
K. Schnaitter and N. Polyzotis.
ACM Transactions on Database Systems, vol. 35, no. 1, 2010.
4. “TuG Synopses for Approximate Query Answering”.
J. Spiegel and N. Polyzotis.
ACM Transactions on Database Systems, vol. 34, no. 1, 2009.
5. “Depth Estimation for Ranking Query Optimization”.
K. Schnaitter, J. Spiegel, and N. Polyzotis.
International Journal on Very Large Databases (VLDBJ), vol. 18, no. 2, 2009.
6. “Meshing Streaming Updates with Persistent Data in an Active Data Warehouse”.
N. Polyzotis, S. Skiadopoulos, P. Vassiliadis, A. Simitsis, and N.E. Frantzell.
In IEEE Transactions on Knowledge and Data Engineering (TKDE), vol. 20, no. 7, 2008.
7. “XSketch Synopses for XML Data Graphs”.
N. Polyzotis and M. Garofalakis.
In ACM Transactions on Database Systems, vol. 29, no. 3, 2006.

Articles in Top-Tier Peer Reviewed Conference Proceedings

8. “Optimal Top-K Query Evaluation for Weighted Business Processes”.
D. Deutch, T. Milo, N. Polyzotis, and T. Yam.
In Proceedings of the VLDB Endowment (PVLDB), 3(1), 2010.
9. “An automated, yet interactive and portable DB designer”.
I. Alagiannis, D. Dash, K. Schnaitter, A. Ailamaki, N. Polyzotis.
In Proceedings of ACM SIGMOD International Conf. on Management of Data
(Demonstrations Track), pp. 1183-1186, 2010.
10. “Index Interactions in Physical Design Tuning: Modeling, Analysis, and Applications”.
K. Schnaitter, N. Polyzotis, and L. Getoor.
In Proceedings of the VLDB Endowment (PVLDB), 2(1), pp. 1234-1245, 2009.
11. “Autocompletion for Mashups”.
O. Greenshpan, T. Milo, and N. Polyzotis.
In Proceedings of the VLDB Endowment (PVLDB), 2(1), pp. 538-549, 2009.
12. “A Scalable, Predictable Join Operator for Highly Concurrent Data Warehouses”.
G. Candea, N. Polyzotis, and R. Vingralek.
In Proceedings of the VLDB Endowment (PVLDB), 2(1), pp. 277-288, 2009.

13. "Robust and Efficient Algorithms for Rank Join Evaluation".
J. Finger and N. Polyzotis
In Proceedings of ACM SIGMOD International Conf. on Management of Data, pp. 415-428, 2009.
14. "MatchUp: Autocompletion for Mashups".
S. Abiteboul, O. Greenshpan, T. Milo, and N. Polyzotis.
In Proceedings of IEEE International Conf. on Data Engineering (ICDE) (Demonstrations Track), pp. 1479-1482, 2009.
15. "Sketch-Based Summarization of Ordered XML Streams".
V. Mayorga and N. Polyzotis.
In Proceedings of IEEE International Conf. on Data Engineering (ICDE), pp. 541-552, 2009.
16. "Evaluating rank joins with optimal cost".
K. Schnaitter and N. Polyzotis.
In Proceedings of ACM Symposium on Principles of Database Systems (PODS), pp. 43-52, 2008.
17. "XML Processing in DHT Networks".
S. Abiteboul, I. Manolescu, N. Polyzotis, N. Preda, C. Sun.
In Proceedings of IEEE International Conf. on Data Engineering (ICDE), pp. 606-615, 2008.
18. "The Data Ring: Community Content Sharing".
S. Abiteboul and N. Polyzotis.
In Proceedings of Conference on Innovative Data systems Research (CIDR), 2007.
19. "Supporting Streaming Updates in an Active Data Warehouse".
N. Polyzotis, S. Skiadopoulos, P. Vassiliadis, A. Simitsis, and N.E. Frantzell.
In Proceedings of IEEE International Conf. on Data Engineering (ICDE), pp. 476-485, 2007.
20. "Depth Estimation for Ranking Query Optimization".
K. Schnaitter, J. Spiegel, N. Polyzotis.
In Proceedings of International Conf. on Very Large Data Bases (VLDB), pp. 902-913, 2007.
21. "Graph-Based Synopses for Relational Selectivity Estimation".
J. Spiegel and N. Polyzotis.
In Proceedings of ACM SIGMOD International Conf. on Management of Data, pp. 205-216, 2006.
22. "AQAX: A System for Approximate XML Query Answers".
J. Spiegel, E. Pontikakis, S. Budalakoti, and N. Polyzotis.
In Proceedings of International Conf. on Very Large Data Bases (VLDB) (Demonstrations Track), pp. 1159-1162, 2006.
23. "COLT: Continuous On-Line Database Tuning".
K. Schnaitter, S. Abiteboul, T. Milo, and N. Polyzotis.
In Proceedings of ACM SIGMOD International Conf. on Management of Data (Demonstrations Track), pp. 793-795, 2006.

24. "XCluster Synopses for Structured XML Content".
N. Polyzotis and M. Garofalakis.
In Proceedings of IEEE International Conf. on Data Engineering (ICDE), pp. 63-72, 2006.
25. "Approximate XML Query Answers".
N. Polyzotis, M. Garofalakis, and Y. Ioannidis.
In Proceedings of ACM SIGMOD International Conf. on Management of Data, pp. 263-274, 2004
26. "Selectivity Estimation for XML Twigs".
N. Polyzotis, M. Garofalakis, and Y. Ioannidis.
In Proceedings of IEEE International Conf. on Data Engineering (ICDE), pp. 264-275, 2004.
27. "Speculative Query Processing".
N. Polyzotis and Y. Ioannidis.
In Proceedings of International Conf. on Innovative Data Systems Research (CIDR), 2003.
28. "Structure and Value Synopses for XML Data Graphs".
N. Polyzotis and M. Garofalakis.
In Proceedings of International Conf. on Very Large Data Bases (VLDB), pp. 466-477, 2002.
29. "Statistical Synopses for Graph-Structured XML Databases".
N. Polyzotis and M. Garofalakis.
In Proceedings of ACM SIGMOD Intl. Conf. on Management of Data, pp. 358-369, 2002.

Articles in other Peer Reviewed Conference and Workshop Proceedings

30. "Fusing Data Management Services with File Systems".
S. Brandt, C. Maltzhan, W.-C. Tan, and N. Polyzotis.
In Proceedings of the 4th Petascale Data Storage Workshop, Supercomputing '09.
31. "Query Recommendations for Interactive Database Exploration".
G. Chatzopoulou, M. Eirinaki, and N. Polyzotis.
In Proceedings of International Conf. on Scientific and Statistical Database Management (SSDBM), pp. 3-18, 2009.
32. "A Benchmark for On-line Index Selection".
K. Schnaitter and N. Polyzotis.
In Proceedings of IEEE International Conf. on Data Engineering (ICDE) (Workshop on Self-Managing Database Systems), pp. 1701-1708, 2009.
33. "On-Line Index Selection for Shifting Workloads".
Karl Schnaitter, Serge Abiteboul, Tova Milo, Neoklis Polyzotis.
In Proceedings of IEEE International Conf. on Data Engineering (ICDE) Workshops, pp. 459-468, 2007.
34. "Tuple Graph Synopses for Relational Data Sets".
J. Spiegel and N. Polyzotis
In 5th Hellenic Data Management Symposium, 2006.

35. "Selectivity-Based Partitioning: A Divide and Union Technique for Effective Query Processing".
N. Polyzotis.
In Proceedings of ACM International Conf. on Knowledge and Information Management (CIKM), pp. 720-727, 2005.
36. "Searching a File System using Inferred Semantic Links".
D. Bhagwat and N. Polyzotis.
In Proceedings of ACM Hypertext, pp. 85-87, 2005.
37. "Approximate Answers for XML Queries with Range Predicates".
N. Polyzotis and M. Garofalakis.
In Proceedings of 4th Hellenic Data Management Symposium, 2005.
38. "Fractional XSketch Synopses for XML Databases".
N. Drukh, N. Polyzotis, M. Garofalakis, and Y. Matias.
In Proceedings International XML Database Symposium (XSym), pp. 189-203, 2004.
39. "Approximate Twig-Query Processing".
N. Polyzotis, M. Garofalakis, and Y. Ioannidis.
In 3rd Hellenic Data Management Symposium, 2004.
40. "XSketch Synopses for XML".
N. Polyzotis and M. Garofalakis.
In 1st Hellenic Data Management Symposium, 2002.