

# Semih Salihoglu

---

CONTACT INFORMATION	200 University Ave W, Waterloo, ON N2L 3G1 semih.salihoglu@uwaterloo.ca
RESEARCH INTERESTS	Graph Databases, Data Management and Processing Systems, Distributed Systems  My research focuses on new theories, algorithms, and systems for processing and managing data. Specifically, I am interested in Graph databases and Pregel-like systems that process data in the form of a graph.
EDUCATION	<b>Stanford University</b> , Stanford, CA Ph.D., Computer Science 2009-2015 <ul style="list-style-type: none"><li>• Advisor: Jennifer Widom</li><li>• Thesis Title: Massive-scale Processing of Record-Oriented and Graph Data</li></ul> <b>Yale University</b> , New Haven, CT B.S., Computer Science and Economics, 2002-2006 <ul style="list-style-type: none"><li>• Summa cum laude, Phi Beta Kappa</li></ul>
WORK EXPERIENCE	<b>University of Waterloo. Assistant Professor, Waterloo, ON</b> 2016-present <b>Google, Inc. Software Engineer, New York, NY</b> 2006-2009
PUBLICATIONS	<b>Conference Papers</b> <ol style="list-style-type: none"><li>1. Spectral Measures of Distortion for Change Detection in Dynamic Graphs, Aleardi L. C., Salihoglu S., Singh G., Ovsjanikov M., Complex Networks, December 2018</li><li>2. Distributed Evaluation of Subgraph Queries Using Worstcase Optimal Low-Memory Dataflows, Ammar K., McSherry F., Salihoglu S., Proc. International Conference on Very Large Data Bases (VLDB), August 2018</li><li>3. The Ubiquity of Large Graphs and Surprising Challenges of Graph Processing, Sahu S., Mhedhbi A., Salihoglu S., Lin J., Ozsü T., Proc. International Conference on Very Large Data Bases (VLDB), August 2018 (<i>Best Paper Award</i>)</li><li>4. Workload-Aware CPU Performance Scaling for Transactional Database Systems, Korkmaz M., Karsten M., Salem K., Salihoglu S., ACM SIGMOD International Conference on Management of Data (SIGMOD), June 2018</li><li>5. GYM: A Multiround Join Algorithm In MapReduce, Afrati F., Joglekar M., Re C., Salihoglu S., Ullman J., Proc. International Conference on Database Theory (ICDT), March 2017</li><li>6. Optimizing Graph Algorithms on Pregel-like Systems, Salihoglu, S., Widom, J., Proc. International Conference on Very Large Data Bases (VLDB), September 2014</li><li>7. Anchor Points Algorithms for Hamming and Edit Distance Afrati F., Das Sarma, A., Rajaraman, A., Rule, P., Salihoglu, S., Ullman, J., Proc. International Conference on Database Theory (ICDT), March 2014</li><li>8. Simplifying Scalable Graph Processing with a Domain-Specific Language Hong., S., Salihoglu., S., Widom., J., Olukotun., K. Proc. International Symposium on Code Generation and Optimization (CGO), February 2014</li></ol>

9. Upper and Lower Bounds on the Cost of a MapReduce Computation Afrati, F., Das Sarma, A., Salihoglu, S., Ullman J., Proc. International Conference on Very Large Data Bases (VLDB), August 2013
10. GPS: A Graph Processing System Salihoglu, S., Widom, J. Proc. International Conference on Scientific and Statistical Database Management (SSDBM), July 2013 (*Best Paper Runner-up Award*)
11. Provenance-Based Refresh in Data-Oriented Workflows Ikeda, R. Salihoglu, S., Widom, J. Proc. Conference on Information and Knowledge Management (CIKM), October 2011

### Journal Papers

1. Algorithmic Aspects of Parallel Data Processing Koutris P., Salihoglu S., Suciu D., Foundations and Trends in Databases, Volume 8, Issue 4, February 2018 (*135 page-long survey paper printed in book format*)
2. Combining Vertex-centric Graph Processing with SPARQL for Large-scale RDF Data Analytics Abdelaziz I., Al-Harbi R., Salihoglu S., and Kalnis P., IEEE Transactions on Parallel and Distributed Systems, June 2017

### Workshop Papers

1. Help: High-level Primitives for Large-Scale Graph Processing, Salihoglu, S., Widom., J., Graph Data-management Experiences and Systems Workshop (GRADES), June 2014
2. Towards an Understanding of the Limits of MapReduce Computation Afrati, F., Das Sarma, A., Salihoglu, S., Ullman J., Cloud Futures Workshop, Berkeley, CA, May 2012

### Demonstration Papers

1. Graphflow: An Active Graph Database, Proc. ACM SIGMOD International Conference on Management of Data (SIGMOD) (Demonstration Track), May 2017 Kankanamge C., Sahu S., Mhedhbi A., Chen J., Salihoglu S.
2. SPARTex: A Vertex-Centric Framework for RDF Data Analytics Abdelaziz, I. and Harbi, R. and Salihoglu, S. and Kalnis, P. and Mamoulis, Proc. International Conference on Very Large Data Bases (VLDB), September 2015
3. Graft: A Debugging Tool For Apache Graph Salihoglu, S. and Shin, J. and Khanna, V. and Truong, B. and Widom, J., Proc. ACM SIGMOD International Conference on Management of Data (SIGMOD), June 2015
4. Provenance-Based Debugging and Drill-Down in Data-Oriented Workflows Ikeda, R., Cho, J., Fang, C., Salihoglu, S., Torikai, S., Widom, J. Proc. International Conference on Data Engineering (ICDE), Washington, DC, April 2012

### AWARDS AND HONORS

- Best Paper Award, VLDB 2018
- Distinguished PC Member, SIGMOD 2018
- Excellent Presentation Award, VLDB 2014
- Best Paper Runner-up Award, SSDBM 2013

PROFESSIONAL  
ACTIVITIES

- Tutorial on Algorithmic Aspects of Parallel Query Processing, Koutris P., Salihoglu S., Suciu D., ACM SIGMOD International Conference on Management of Data (SIGMOD), June 2018
- Tutorial on Designing Good Algorithms for MapReduce and Beyond, Afrati F. N., Balazinska M., Das Sarma A., Howe B., Salihoglu S., Ullman J., ACM Symposium on Cloud Computing (SOCC), October 2012
- Sponsorship co-chair, SIGMOD 2018, 2019
- Proceedings co-chair, SIGMOD 2017
- PC Memberships:
  - SIGMOD 2017, 2018, 2019
  - VLDB 2017, 2019
  - ICDE 2017, 2018
  - Workshop on Graph Algorithms Building Blocks (GABBS), 2017
  - Workshop on Algorithms for MapReduce and Beyond, 2014, 2015, 2016, 2017
  - High Performance Graph Data Management and Processing Workshop (HPGDMP), 2016
- Reviewer:
  - VLDB Journal, 2017
  - ACM Transactions on Database Systems (TODS), 2018
  - IEEE Transactions on Parallel and Distributed Systems (TPDS), 2017, 2018
  - IEEE Transactions on Big Data, 2017
  - IEEE Transactions on Knowledge and Data Engineering (TKDE) 2013, 2017
  - Communications of the ACM 2015
  - USENIX ATC 2017, EDBT 2014, ICDE 2012 (external)

TEACHING

- CS 848: Graph Data Management, Fall 2018, University of Waterloo
- CS 848: Models and Applications of Distributed Data Processing Systems, Fall 2017, University of Waterloo
- CS 341: Algorithms, Summer 2016, Winter 2017, Winter 2018, University of Waterloo

SERVICE AT  
UNIVERSITY OF  
WATERLOO

- Reader on Yan Zhang's MMath thesis (advised by Tamer Ozsu)
- Reader on Vivi Ma's MMath thesis (advised by Grant Weddell)
- Grad Committee Member (2016-17 and 2017-18)
- UROC instructor (2017)

RECENT TALKS

- Graph Query Processing using Worst-case Optimal Join Algorithms and Differential Computation, Shonan Workshop on Graph Database Systems, Shonan, Japan, July 2018
- Applications of New Join and Spectral Algorithms in Dynamic Graph Processing, Ecole Polytechnique, Paris, France, March 2017
- Surprises in a Survey of the Graphs, Computations, and Software Used in Practice, Workshop on Big Graph Analysis Systems, Copenhagen, Denmark, August 2017
- Graphflow Graph Database, SAP Hana Tech Days, Munich, Germany, August 2017

CURRENT  
GRADUATE  
SUPERVISION

- Amine Mhedhbi, PhD (2016-present) (w/ Tamer Ozsu)
- Siddhartha Sahu, PhD (2016-present) (w/ Tamer Ozsu)
- Kaleb Alway, MMath (2017-present) (w/ Eric Blais)
- Jeremy Chen, MMath (2017-present)
- Shahid Khaliq, MMath (2017-present)
- Arman Naeimian, MMath (2017-present) (w/ Mei Nagappan)
- Azin Nazari, MMath (2017-present) (w/ Ian Munro)

PAST GRADUATE  
SUPERVISION

- Chathura Kankanamge, MMath, 2018. Title: *Multiple Continuous Subgraph Query Optimization Using Delta Subgraph Queries*