

## Semih Salihoglu

---

CONTACT INFORMATION	200 University Ave W Waterloo, ON N2L 3G1	semih.salihoglu@uwaterloo.ca
RESEARCH INTERESTS	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> <b>Google, Inc. Software Engineer, New York, NY</b>	2016-present 2006-2009
AWARDS AND HONORS	<ul style="list-style-type: none"><li>• VMWare Fellowship in Honor of Pratap Subrahmanyam, 2014-2015</li><li>• Stanford Computer Science Department Google Graduate Fellowship, 2013-2014</li><li>• Excellent Presentation Award, VLDB 2014</li><li>• Best Paper Runner-up Award, SSDBM 2013</li><li>• Yale Science and Engineering Senior Highest Scholarship Award, 2006</li><li>• Yale Computer Science Department Undergraduate Award, 2006</li></ul>	
PROFESSIONAL ACTIVITIES	<ul style="list-style-type: none"><li>• 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 (to appear)</li><li>• 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</li><li>• Sponsorship co-chair, SIGMOD 2018</li><li>• Proceedings co-chair, SIGMOD 2017</li><li>• PC Member, SIGMOD 2017, 2018</li><li>• PC Member, ICDE 2017, 2018</li><li>• PC Member, VLDB 2017</li><li>• Reviewer for VLDB Journal, 2017</li><li>• Reviewer for ACM Transactions on Database Systems, 2018</li><li>• Reviewer for IEEE Transactions on Parallel and Distributed Systems (TPDS), 2017</li><li>• Reviewer for IEEE Transactions on Big Data, 2017</li><li>• External Reviewer for USENIX Annual Technical Conference (ATC), 2017</li></ul>	

- PC Member for Workshop on Graph Algorithms Building Blocks (GABBS), 2017
- PC Member for Workshop on Algorithms for MapReduce and Beyond, EDBT/ICDT, 2014, 2015, 2016, 2017
- PC Member for High Performance Graph Data Management and Processing Workshop (HPGDMP), 2016
- Reviewer for IEEE TKDE 2013, 2017, CACM 2015
- External reviewer for EDBT 2014, ICDE 2012

#### TEACHING

- 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
- CS 161: Design and Analysis of Algorithms, Summer 2014, Stanford University

#### PUBLICATIONS

##### Conferences

1. 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), Rio de Janeiro, August 2018 (to appear)
2. The Ubiquity of Large Graphs and Surprising Challenges of Graph Processing, Sahu S., Mhedhbi A., Salihoglu S., Lin J., zsu T., Proc. International Conference on Very Large Data Bases (VLDB), Rio de Janeiro, August 2018 (to appear)
3. GYM: A Multiround Join Algorithm In MapReduce, Afrati F., Joglekar M., Re C., Salihoglu S., Ullman J., Proc. International Conference on Database Theory (ICDT), Venice, Italy, March 2016
4. Optimizing Graph Algorithms on Pregel-like Systems, Salihoglu, S., Widom, J., Proc. International Conference on Very Large Data Bases (VLDB), Hangzhou, China, September 2014
5. 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), Athens, Greece, March 2014
6. 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), Orlando, FL, February 2014
7. 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), Trento, Italy, August 2013
8. GPS: A Graph Processing System Salihoglu, S., Widom, J. Proc. International Conference on Scientific and Statistical Database Management (SSDBM), July 2013
9. Provenance-Based Refresh in Data-Oriented Workflows Ikeda, R. Salihoglu, S., Widom, J. Proc. Conference on Information and Knowledge Management (CIKM), Glasgow, Scotland, October 2011

##### Journals

1. Algorithmic Aspects of Parallel Data Processing Koutris P., Salihoglu S., Suciu D., Foundations and Trends in Databases, Volume 8, Issue 4, February 2018

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

### **Workshops**

1. HeLP: High-level Primitives for Large-Scale Graph Processing, Salihoglu, S., Widom., J., Graph Data-management Experiences and Systems Workshop (GRADES), Snowbird, Utah, 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
3. Interdomain Routing As Social Choice Dakdouk, R., Salihoglu, S., Wang, H., Xie, Yang Y., Distributed Computing Systems Workshops (ICDCS) , Lisbon, Portugal, July 2006

### **Demonstrations**

1. Graphflow: An Active Graph Database, Proc. ACM SIGMOD International Conference on Management of Data (SIGMOD) (Demonstration Track), May 2017 Kankanamge C., Sahu S., Mhedbhi 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