

Shumo Chu

Curriculum Vitae

Box 352350

Seattle, WA 98195

☎ (206) 422 8431

✉ chushumo@cs.washington.edu

🌐 shumochu.com

Education

- 2013–Present **Ph.D. Student**, *Paul G. Allen School of Computer Science & Engineering, University of Washington, WA*, advised by Prof. Dan Suciu.
- 2010–2012 **Master of Computer Science**, *Nanyang Technological University, Singapore*.
- 2005–2009 **Bachelor of Engineering**, *Wuhan University, China*.

Research Projects

- 2015–Present **Cosette – An Automated SQL Solver**.
Cosette is the first automated reasoning tool for checking the semantic equivalences of SQL queries. It leverages both an interactive theorem prover and a constraint solver. Given a pair of SQL queries, Cosette can either find mechanized equivalence proofs or show that they are not equivalent by providing a counterexample. Cosette covers major SQL features and is efficient in practice. (<http://cosette.cs.washington.edu/>)
- 2013–Present **Myria – Distributed Database System as a Service**.
Myria is a shared nothing distributed database system developed in University of Washington. I worked on developing a novel multiway join processing subsystem for Myria, which achieves 10x better performance by reducing the number of intermediate tuples generated and shuffled. (<https://github.com/uwescience/myria>)

Awards

- 2017 SIGMOD Best Demo Award
- 2012 “Best of KDD 2011” citation in the ACM TKDD

Industry Experience

- June–Sep. **Research Intern**, MICROSOFT RESEARCH, Redmond, WA.
 - 2015 Design and implement a new sampling framework used for analytic query cardinality estimation (mentor: Christian Konig, Bolin Ding)
- March–June **Software Engineering Intern**, GOOGLE, Kirkland, WA.
 - 2015 Worked with Spanner team on distributed database query optimization.

Publications

- 2018 **Shumo Chu**, Brendan Murphy, Jared Roesch, Alvin Cheung, Dan Suciu. Axiomatic Foundations and Algorithms for Deciding Semantic Equivalences of SQL Queries. *International Conference on Very Large Data Bases (VLDB)*.

- 2017 **Shumo Chu**, Daniel Li, Chenglong Wang, Alvin Cheung and Dan Suciu. Demonstration of the Cosette Automated SQL Prover. *ACM SIGMOD conference*. **Best Demo Award**.
- 2017 **Shumo Chu**, Konstantin Weitz, Alvin Cheung, Dan Suciu. HoTTSQL: Proving Query Rewrites with Univalent SQL Semantics. *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.
- 2017 **Shumo Chu**, Chenglong Wang, Konstantin Weitz, Alvin Cheung. Cosette: A Automated Prover for SQL. *The biennial Conference on Innovative Data Systems Research (CIDR)*.
- 2015 **Shumo Chu**, Magdalena Balazinska and Dan Suciu. From Theory to Practice: Efficient Join Query Evaluation in a Parallel Database System. *ACM SIGMOD Conference*.
- 2014 with the Myria Team. Demo of the Myria Big Data Management Service. *ACM SIGMOD conference*.
- 2012 **Shumo Chu** and James Cheng. Triangle Listing in Massive Network. *ACM Transaction on Knowledge Discovery from Data*.
- 2012 James Cheng, Linhong Zhu, Yiping Ke and **Shumo Chu**. Fast Algorithms for Maximal Clique Enumeration with Bounded Memory. *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*.
- 2012 James Cheng, Yiping Ke **Shumo Chu** and Carter Cheng. Efficient Processing of Distance Queries in Large Graphs: A Vertex Cover Approach. *ACM SIGMOD Conference*.
- 2011 **Shumo Chu** and James Cheng. Triangle Listing in Massive Networks and Its Applications. *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, **Among bests of SIGKDD 2011**.
- 2011 James Cheng, Yiping Ke, **Shumo Chu** and Tamer Öszu. Core Decomposition in Massive Networks. *IEEE International Conference on Data Engineering (ICDE)*.

Services

- 2018 Reviewer of Knowledge and Information System, IEEE Transactions on Knowledge and Data Engineering
- 2017 PLDI Artifact Evaluation Committee, ACM CIKM PC

Referees

Dan Suciu, *Advisor*, Professor, University of Washington.
suciu@cs.washington.edu

Zachary Tatlock, *Thesis Committee*, Assistant Professor, University of Washington.
ztatlock@cs.washington.edu

Magdalena Balazinska, *Collaborator*, Professor, University of Washington.
magda@cs.washington.edu

James Cheng, *Master Advisor*, Assistant Professor, The Chinese University of Hong Kong.

jcheng@cse.cuhk.edu.hk

Christian Konig, *Mentor*, Senior Researcher, DMX Group, Microsoft Research.

chrisko@microsoft.com