## Dr. Zuyu ZHANG

 $zuyu@cs.wisc.edu \\ http://www.cs.wisc.edu/~zuyu$ 

EDUCATION	<ul> <li>Ph.D., Database Systems</li> <li>M.S., Computer Sciences</li> <li>University of Wisconsin-Madison, Madison, WI, US</li> <li>Advisor: Prof. Jignesh Patel</li> <li>Thesis: Towards Efficient Data Processing Methods For In-Memory Arch</li> </ul>	2019 2013 hitectures.
	<b>B.Eng., Computer Science and Technology</b> Harbin Institute of Technology (HIT), Harbin, China	2010
TECHNICAL SKILLS	Prog. Lang.: C/C++, Python, Java, and Rust. Tools: AWS, Docker, CMake, Protobuf, gTest, gRPC, GIT, and LATEX.	
EXPERIENCE	<ul> <li>Sr. Software Engineer, Celonis Inc Jul 2022-Sep 2023</li> <li>Prototyped a new cloud-native Process Mining engine based on StarRocks.</li> <li>Added specialized operators for graph functions in Process Query Language.</li> </ul>	
	<ul> <li>Software Engineer II, Amazon Web Services Sep 20</li> <li>Worked on Aurora MySQL (AMS), a cloud-native managed datal</li> <li>Launched AMS 8.0 as a team.</li> <li>Improved AMS 5.7 functionality, availability, reliability, and security</li> </ul>	019-May 2022 base service. ty.
	<ul> <li>Research Assistant, UW-Madison Sep 2017-Sep 2019</li> <li>Worked on Quickstep to support TPC-H queries in C++.</li> <li>Prototyped RecStep, the first efficient datalog engine on DBMS, including the bit-matrix approach for Transitive Closure (TC) and Same Generation (SG).</li> <li>Worked towards a unified data engine for both SQL and graph analytics in Rust.</li> </ul>	
	<ul> <li>ACM SIGMOD Programming Contest</li> <li>Built an in-memory analytical query engine from scratch w/ J. Zhu</li> <li>Proposed the join-aggregate operator fusion, and the index-based l</li> <li>Winner among 30+ teams from 10+ countries; 42% faster than the</li> </ul>	Feb-Apr 2018 u in C++. inear-join. he runner-up.
	<ul> <li>Software Engineering Intern, Google</li> <li>Extended an interactive reporting service towards multi-tenant in</li> </ul>	Oct-Dec 2016 $C++$ .
	Software Engineer Intern, Pivotal       Jun 2         • Research and develop distributed analytical SQL query processing	2015-Jul 2016 techniques.
	<ul> <li>Software Engineer Intern, Twitter</li> <li>Prototyped Heron Scheduler on Apache Mesos with automatic faile</li> <li>Demonstrated the scheduler running on a Mesos cluster of 150 nod</li> </ul>	/lay-Aug 2014 over. les.
SELECTED PUBLICATIONS	Zhiwei Fan, Jianqiao Zhu, <b>Zuyu Zhang</b> , Paraschos Koutris, Aws Albarghouthi, Jig- nesh M. Patel, "Scaling-up In-Memory Datalog Processing: Observations and Tech- niques", <i>Proc.</i> <b>VLDB</b> <i>Endow.</i> 12, 6 (February 2019), 695-708.	

**Zuyu Zhang**, Harshad Deshmukh, Jignesh M. Patel, "Data Partitioning for In-Memory Systems: Myths, Challenges, and Opportunities", 9th Biennial Conference on Innovative Data Systems Research (CIDR), 2019.

Jignesh M. Patel, Harshad Deshmukh, Jianqiao Zhu, Navneet Potti, **Zuyu Zhang**, Marc Spehlmann, Hakan Memisoglu, Saket Saurabh, "Quickstep: A Data Platform Based on the Scaling-Up Approach", *Proc.* **VLDB** *Endow.* 11, 6 (February 2018), 663-676.

**INTERESTS** Sailing, clarinet, and travel (10 countries).