

Matthew Perron

mperron@mit.edu

+1-(802)-238-3141

matthew-perron.com

RESEARCH AREA	Cloud Data Systems My research makes data systems easier for non-expert users to deploy and use to answer their most important questions. My work focuses on how analytical database systems can use elasticity to meet user performance goals in dynamic workloads and environments while maintaining low cost.
EDUCATION	Ph.D. Computer Science , Massachusetts Institute of Technology EECS (exp.) 2024 M.S. Computer Science , Carnegie Mellon University, CSD 2017 B.S. Computer Science , Rochester Institute of Technology 2013 Minors: Japanese Language/Japanese Language and Culture Study Abroad: Sophia University, Tokyo, Japan 2010-2011
EXPERIENCE	Research Assistant, MIT CSAIL 2017–2023 <i>Advisor:</i> Samuel Madden <i>Coauthors:</i> David DeWitt, Raul Castro Fernandez, Michael Cafarella, Tim Kraska, Michael Stonebraker <i>Research Topics:</i> Cloud Analytical Databases, Query Reoptimization, Cloud Database Benchmarking <i>Publication Venues:</i> SIGMOD, VLDB, ICDE Intern, Amazon Redshift Summer 2020 and 2021 Supervised by Ippokratis Pandis <ul style="list-style-type: none">• Researched techniques to extending Redshift with elastic stateless compute to reduce query latency.• Received patent for this work. Research Intern, Microsoft Research, DMX Group Summer 2018 Supervised by Srikanth Kandula, Surajit Chaudhuri <ul style="list-style-type: none">• Researched techniques for cardinality estimation using machine learning.
HONORS & AWARDS	NSF GRFP Honorable Mention 2017 Irwin Mark Jacobs (1957) and Joan Klein Jacobs Presidential Fellow 2017-2018
TEACHING EXPERIENCE	Teaching Assistant, MIT 6.S080: Software Systems for Data Science Fall 2019 <ul style="list-style-type: none">• Ran office hours once weekly for a class of 74 students.• Designed and graded half of course assignments• Gave a lecture Guest Lecture, MIT 6.887: Machine Learning For Systems Fall 2021 <ul style="list-style-type: none">• Gave a guest lecture on Spark execution to a class of 24 students. Teaching Assistant, MIT Brave Behind Bars Summer 2023 <ul style="list-style-type: none">• Ran twice weekly sessions with small groups of incarcerated students, reinforcing course material.
OTHER ACTIVITIES	President, MIT Rowing Club 2020-2022 <ul style="list-style-type: none">• Responsible for the period during and after COVID pandemic, growing the club to 100 members.• Planned season details, arranged coaches, and coach payments.• Recruited club officers, organized regular officer meetings. Member, CSAIL Researcher Council 2021-2023 <ul style="list-style-type: none">• Participated in regular meetings with the lab director.

- Brought concerns from students to the attention of lab administration.

Mentor, MIT Graduate Application Assistance Program 2020

- Advised two students from underrepresented minorities in revising PhD application materials over several meetings.
- Both students now enrolled in PhD programs.

New Member Manager, MIT Rowing Club 2022-2023

- Responsible for communicating with prospective club members.

TALKS	Cockroach Labs	January 2020
	University of Chicago ChiData Group	April 2020
	Databricks	June 2020
	IBM	Sept 2020

PUBLICATIONS [1] **Matthew Perron**, Michael Cafarella, Raul Castro Fernandez, David DeWitt, and Samuel Madden. Cackle: Analytical Workload Cost and Performance Stability With Elastic Pools (To Appear) . *Proc. ACM Manag. Data*, 1(1), December 2024

[2] **Matthew Perron**, Raul Castro Fernandez, David DeWitt, and Samuel Madden. Starling: A Scalable Query Engine on Cloud Functions. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, SIGMOD '20, page 131–141, New York, NY, USA, 2020. Association for Computing Machinery

[3] Junjay Tan, Thanaa Ghanem, **Matthew Perron**, Xiangyao Yu, Michael Stonebraker, David DeWitt, Marco Serafini, Ashraf Aboulnaga, and Tim Kraska. Choosing a Cloud DBMS: Architectures and Trade-offs. *Proceedings of the VLDB Endowment*, 12(12):2170–2182, 2019

[4] **Matthew Perron**, Zeyuan Shang, Tim Kraska, and Michael Stonebraker. How I Learned to Stop Worrying and Love Re-optimization. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 1758–1761. IEEE, 2019

[5] Andrew Pavlo, Gustavo Angulo, Joy Arulraj, Haibin Lin, Jiexi Lin, Lin Ma, Prashanth Menon, Todd C Mowry, **Matthew Perron**, Ian Quah, et al. Self-Driving Database Management Systems. In *CIDR*, volume 4, page 1, 2017

[6] Joy Arulraj, **Matthew Perron**, and Andrew Pavlo. Write-Behind Logging. *Proceedings of the VLDB Endowment*, 10(4):337–348, 2016

PATENTS [7] Ippokratis Pandis and **Matthew James Perron**. Selecting Between Hydration-Based Scanning and Stateless Scale-out Scanning to Improve Query Performance, June 22 2023. US Patent App. 18/171,245