

Markos Markakis

Email: markakis@mit.edu **Phone:** (617) 685-9808

Research Interests

I am interested in tackling the efficient management of big data by designing novel high-performance data systems. I consider both machine learning techniques and a fine-grained knowledge of new hardware trends as particularly useful tools to that end.

Education

Massachusetts Institute of Technology Cambridge, MA
 PhD Candidate in Electrical Engineering and Computer Science Jun 2022 – Present
 • Advisor: Prof. Tim Kraska ([Data Systems Group](#))
 • GPA: 5.00/5.00

Massachusetts Institute of Technology Cambridge, MA
 Master of Science in Electrical Engineering and Computer Science Sep 2020 – May 2022
 • Thesis: Rethinking Update-in-Place Key-Value Stores for Modern Storage
 • Advisor: Prof. Tim Kraska ([Data Systems Group](#))
 • GPA: 5.00/5.00

Princeton University Princeton, NJ
 BSE Summa Cum Laude in Electrical Engineering Sep 2016 – Jun 2020
 • Thesis: Challenges and Opportunities in Heterogeneous Parallelism
 • Advisor: Prof. Margaret Martonosi ([MRM Research Group](#))
 • GPA: 3.99/4.00

Research Experience

Massachusetts Institute of Technology Cambridge, MA
 Graduate Research Assistant, [Data Systems Group](#) Sep 2020 – Present
 • Working on utilizing machine learning to accelerate database systems.
 • Collaborated with Microsoft Research, Intel and Amazon as part of MIT's [DSAIL](#) effort.

Amazon Web Services Boston, MA
 Applied Scientist Intern May 2023 – Aug 2023
 • Interned with the Learned Systems Group within Amazon Redshift.
 • Profiled recurring customer workloads and developed solutions to accelerate them.

Intel Corporation Munich, Germany
 Graduate Research Intern Jun 2021 – Aug 2021
 • Interned remotely with the US-based Intel Cloud Enterprise Solutions Group.
 • Explored solutions for leveraging persistent memory technology for key-value stores.

Princeton University Princeton, NJ
 Undergraduate Research Assistant, [MRM Research Group](#) Feb 2019 – Jun 2020
 • Contributed to [DECADES](#) by co-developing the formal verification infrastructure.
 • Worked on the PerpLE project, a novel approach for improving memory consistency testing.

Publications

From Logs to Causal Inference: Diagnosing Large Systems
 Markos Markakis, Brit Youngmann, Trinity Gao, Ziyu Zhang, Rana Shahout, Peter Baile Chen, Chunwei Liu, Ibrahim Sabek, Michael Cafarella.
Accepted, Proceedings of the VLDB Endowment 18, 2024.

Blueprinting the Cloud: Unifying and Automatically Optimizing Cloud Data Infrastructures with BRAD
 Geoffrey X. Yu, Ziniu Wu, Ferdi Kossmann, Tianyu Li, Markos Markakis, Amadou Ngom, Samuel Madden, and Tim Kraska
Proceedings of the VLDB Endowment 17 (11), 3629 - 3643, 2024.

Press ECCS to Doubt (Your Causal Graph) - Best Paper Award

Markos Markakis, Ziyu Zhang, Rana Shahout, Trinity Gao, Chunwei Liu, Ibrahim Sabek, Michael Cafarella.

2024 Workshop on Governance, Understanding and Integration of Data for Effective and Responsible AI (GUIDE-AI), 2024.

Sawmill: From Logs to Causal Diagnosis of Large Systems

Markos Markakis, An Bo Chen, Brit Youngmann, Trinity Gao, Ziyu Zhang, Rana Shahout, Peter Baile Chen, Chunwei Liu, Ibrahim Sabek, Michael Cafarella.

2024 ACM SIGMOD/PODS International Conference on Management of Data, 2024.

Check Out the Big Brain on BRAD: Simplifying Cloud Data Processing with Learned Automated Data Meshes

Tim Kraska*, Tianyu Li*, Samuel Madden*, **Markos Markakis***, Amadou Ngom*, Ziniu Wu*, Geoffrey X. Yu*. (* denotes equal contribution)

Proceedings of the VLDB Endowment 16 (11): 3293-3301, 2023.

TreeLine: An Update-In-Place Key-Value Store for Modern Storage

Geoffrey Yu*, **Markos Markakis***, Andreas Kipf*, Per-Åke Larson, Umar Farooq Minhas, Tim Kraska. (* denotes equal contribution)

Proceedings of the VLDB Endowment 16 (1), 99-112, 2022.

PerpLE: Improving the Speed and Effectiveness of Memory Consistency Testing

Themis Melissaris, **Markos Markakis**, Kelly Shaw, Margaret Martonosi.

53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), 2020.

Honors and Awards

Best Paper Award, GUIDE-AI Workshop @ SIGMOD 2024	2024
Paris Kanellakis Fellowship (First-year fellowship, MIT EECS)	2020 – 2021
Calvin Dodd MacCracken Senior Thesis/Project Award	2020
(Most distinctive thesis, Princeton University School of Engineering and Applied Science)	
Charles Ira Young Memorial Tablet and Medal	2020
(Excelled in undergraduate research, Princeton University Dept. of Electrical Engineering)	
Phi Beta Kappa Honor Society	2019
Tau Beta Pi Engineering Honor Society	2018
George B. Wood Legacy Sophomore Prize (Top in class year, Princeton University)	2018
Shapiro Prize for Academic Excellence (Top 3% of class year, Princeton University)	2017, 2018

Teaching Experience

Graduate Teaching Assistant, Massachusetts Institute of Technology	Cambridge, MA
Programming with Data Workshop	Jan 2024
6.S079: Software Systems for Data Science	Feb 2022 – May 2022
Undergraduate Teaching Assistant, Princeton University	Princeton, NJ
ELE 308: Electronic and Photonic Devices	Sep 2019 – Jan 2020
COS 306: Contemporary Logic Design	Sep 2019 – Jan 2020
Undergraduate Computer Science Lab	Feb 2019 – Jun 2019
COS 306: Contemporary Logic Design	Sep 2018 – Jan 2019
COS 217: Introduction to Programming Systems	Sep 2018 – Jan 2019
COS 226: Algorithms and Data Structures	Feb 2018 – Jun 2018

Skills

Programming

- Proficient in: Python, C++, C.
- Familiar with: Verilog, Java, R, MATLAB, LaTeX, Go.

Languages

- Modern Greek (native)
- English (fluent)
- German (advanced)