

Dong Xie

W325 Westgate Building
University Park, PA 16802
+1 (814) 865-9148

dongx@psu.edu
<http://sites.psu.edu/dongx/>

PROFESSIONAL EXPERIENCE

- **Computer Science and Engineering, Pennsylvania State University** University Park, PA
Assistant Professor, August 2020–Present
- **DAMO Academy, Alibaba Group** Sunnyvale, CA
Research Intern, May 2019–August 2019
- **Database Group, Microsoft Research** Redmond, WA
Research Intern, May 2018–August 2018
- **Database Group, Microsoft Research** Sunnyvale, CA
Research Intern, May 2017–August 2017

EDUCATION

- **University of Utah** Salt Lake City, UT
Ph.D. in Computer Science, May 2020
Thesis: *Large Scale Spatial Data Management and Analysis*
- **Shanghai Jiao Tong University** Shanghai, China
B.S.E. in Computer Science, June 2015
Bachelor Thesis: Deterministic Main-Memory Database System

HONORS AND AWARDS

- **2023** Google Research Scholar Award
- **2019** ACM SoCC Best Paper Runner Up
- **2019** ACM SIGMOD Student Travel Grant
- **2018** Microsoft Research Ph.D. Fellow

SELECTED PUBLICATIONS AND ARTIFACTS

CONFERENCE PROCEEDINGS:

- [C9] Zhao Chang, Dong Xie, Feifei Li, Jeff M. Phillips, and Rajeev Balasubramonian. “Efficient and Oblivious Query Processing for Range and kNN Queries (Extended Abstract)”. In: *38th IEEE International Conference on Data Engineering, ICDE 2022, Kuala Lumpur, Malaysia, May 9-12, 2022*. IEEE, 2022, pp. 1487–1488.
- [C8] Zhao Chang, Dong Xie, Sheng Wang, and Feifei Li. “Towards Practical Oblivious Join”. In: *SIGMOD ’22: International Conference on Management of Data*. Ed. by Zachary Ives, Angela Bonifati, and Amr El Abbadi. ACM, 2022, pp. 803–817.
- [C7] Kaisong Huang, Darien Imai, Tianzheng Wang, and Dong Xie. “SSDs Striking Back: The Storage Jungle and Its Implications to Persistent Indexes”. In: *12th Conference on Innovative Data Systems Research, CIDR 2022*. www.cidrdb.org, 2022.
- [C6] Dong Xie, Jeff M. Phillips, Michael Matheny, and Feifei Li. “Spatial Independent Range Sampling”. In: *SIGMOD ’21: International Conference on Management of Data*. Ed. by Guoliang Li, Zhanhuai Li, Stratos Idreos, and Divesh Srivastava. ACM, 2021, pp. 2023–2035.
- [C5] Dong Xie, Badrish Chandramouli, Yinan Li, and Donald Kossmann. “FishStore: Faster Ingestion with Subset Hashing”. In: *Proceedings of the 2019 International Conference on Management of Data, SIGMOD Conference 2019*. Ed. by Peter A. Boncz, Stefan Manegold, Anastasia Ailamaki, Amol Deshpande, and Tim Kraska. ACM, 2019, pp. 1711–1728.

- [C4] Tian Zhang, Dong Xie, Feifei Li, and Ryan Stutsman. “Narrowing the Gap Between Serverless and its State with Storage Functions”. In: *Proceedings of the ACM Symposium on Cloud Computing, SoCC 2019*. ACM, 2019, pp. 1–12.
- [C3] Tao Zhu, Zhuoyue Zhao, Feifei Li, Weining Qian, Aoying Zhou, Dong Xie, Ryan Stutsman, HaiNing Li, and Huiqi Hu. “Solar: Towards a Shared-Everything Database on Distributed Log-Structured Storage”. In: *2018 USENIX Annual Technical Conference, USENIX ATC 2018*. Ed. by Haryadi S. Gunawi and Benjamin Reed. USENIX Association, 2018, pp. 795–807.
- [C2] Dong Xie, Feifei Li, Bin Yao, Gefei Li, Liang Zhou, and Minyi Guo. “Simba: Efficient In-Memory Spatial Analytics”. In: *Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference 2016*. Ed. by Fatma Özcan, Georgia Koutrika, and Sam Madden. ACM, 2016, pp. 1071–1085.
- [C1] Dong Xie, Guanru Li, Bin Yao, Xuan Wei, Xiaokui Xiao, Yunjun Gao, and Minyi Guo. “Practical private shortest path computation based on Oblivious Storage”. In: *32nd IEEE International Conference on Data Engineering, ICDE 2016*. IEEE Computer Society, 2016, pp. 361–372.

JOURNALS, TRANSACTIONS, AND OTHER PERIODICALS:

- [J9] Congying Wang, Nithin Sastry Tellapuri, Sphoorthi Keshannagari, Dylan Zinsley, Zhuoyue Zhao, and Dong Xie. “Approximate Queries over Concurrent Updates”. In: *Proc. VLDB Endow.* 16.12 (2023), pp. 3986–3989.
- [J8] Zhao Chang, Dong Xie, Feifei Li, Jeff M. Phillips, and Rajeev Balasubramonian. “Efficient Oblivious Query Processing for Range and kNN Queries”. In: *IEEE Trans. Knowl. Data Eng.* 34.12 (2022), pp. 5741–5754.
- [J7] Zhuoyue Zhao, Dong Xie, and Feifei Li. “AB-tree: Index for Concurrent Random Sampling and Updates”. In: *Proc. VLDB Endow.* 15.9 (2022), pp. 1835–1847.
- [J6] Michael Matheny, Dong Xie, and Jeff M. Phillips. “Scalable Spatial Scan Statistics for Trajectories”. In: *ACM Trans. Knowl. Discov. Data* 14.6 (2020), 73:1–73:24.
- [J5] Badrish Chandramouli, Dong Xie, Yinan Li, and Donald Kossmann. “FishStore: Fast Ingestion and Indexing of Raw Data”. In: *Proc. VLDB Endow.* 12.12 (2019), pp. 1922–1925.
- [J4] Tao Zhu, Zhuoyue Zhao, Feifei Li, Weining Qian, Aoying Zhou, Dong Xie, Ryan Stutsman, HaiNing Li, and Huiqi Hu. “SolarDB: Toward a Shared-Everything Database on Distributed Log-Structured Storage”. In: *ACM Trans. Storage* 15.2 (2019), 11:1–11:26. DOI: 10.1145/3318158. URL: <https://doi.org/10.1145/3318158>.
- [J3] Dong Xie, Feifei Li, and Jeff M. Phillips. “Distributed Trajectory Similarity Search”. In: *Proc. VLDB Endow.* 10.11 (2017), pp. 1478–1489.
- [J2] Zhao Chang, Dong Xie, and Feifei Li. “Oblivious RAM: A Dissection and Experimental Evaluation”. In: *Proc. VLDB Endow.* 9.12 (2016), pp. 1113–1124.
- [J1] Feifei Li, Ke Yi, Yufei Tao, Bin Yao, Yang Li, Dong Xie, and Min Wang. “Exact and approximate flexible aggregate similarity search”. In: *VLDB J.* 25.3 (2016), pp. 317–338.

SERVICE TO THE PROFESSION

- SIGMOD Program Committee Member, 2022-2024.
- PVLDB Review Board Member, 2022-2023.
- SIGMOD Reproducibility Committee Member, 2021, 2023.
- ICDE Programm Committee Member, 2022.
- VLDB Demonstration Track PC Member, 2021.
- DASFAA Technical PC Member, 2021.
- ACM SIGSPATIAL COVID Workshop PC Member, 2021.
- TODS Reviewer, 2020.
- VLDBJ Reviewer, 2020.
- TKDE Reviewer, 2018-2020.
- VLDB External Reviewer, 2017.

TEACHING

Computer Science and Engineering Department, Pennsylvania State University:

- FA 2023 CMPSC 431W: Database Management Systems 2 Sections Undergrad
- SP 2023 CSE 541: Database Systems I 1 Section Graduate
- SP 2022 CSE 541: Database Systems I 1 Section Graduate
- FA 2021 CSE 597.4: Modern Database System Design 1 Section Graduate
- SP 2021 CMPSC 431W: Database Management Systems 2 Sections Undergrad
- FA 2020 CSE 511: Operating System Design I 1 Section Graduate

GRADUATE STUDENTS

Douglas Rumbaugh, Ph.D, since 2021

Maxwell Norfolk, Ph.D, since 2022

Madhav Venkatesh Deshpande, M.S., graduated 2021