Stephanie Wang

+1 (914)-274-7452
swang@cs.berkeley.edu
stephanie-wang.github.io
stephanie-wang

Research Interests

Distributed Systems, Computer Systems, Programming Languages

Education

2023 (expected) **Ph.D.**, *University of California, Berkeley*, Berkeley, CA.
Advisor: Ion Stoica

2016 **M.Eng.**, *Massachusetts Institute of Technology*, Cambridge, MA.
Advisors: M. Frans Kaashoek and Nickolai Zeldovich
Thesis title: Certifying checksum-based logging in the RapidFSCQ crash-safe filesystem

2015 **B.S.**, *Massachusetts Institute of Technology*, Cambridge, MA.
Computer Science, Mathematics

Publications

preprint **Exoshuffle: Distributed Shuffle as a Library**.
Frank Luan, Stephanie Wang, Samyu Yagati, Sean Kim, Kenneth Lien, Isaac Ong, SangBin Cho, Eric Liang, Ion Stoica.
arxiv.org/abs/2203.05072
Under submission.

Peer-reviewed papers

SoCC'22 **ESCHER: Expressive Scheduling with Ephemeral Resources**.
Romil Bhardwaj, Alexey Tumanov, Stephanie Wang, Richard Liaw, Philipp Moritz, Robert Nishihara, Ion Stoica.
13th ACM Symposium on Cloud Computing (SoCC 2022).

VLDB'22 **Rearchitecting in-memory object stores for low latency**.
Danyang Zhuo, Kaiyuan Zhuang, Zhuohan Li, Siyuan Zhuang, Stephanie Wang, Ang Chen, Ion Stoica.
Proceedings of the Very Large Data Bases Endowment (VLDB 2022).

SIGCOMM'21 **Hoplite: Efficient and Fault-Tolerant Collective Communication for Task-Based Distributed Systems**.
Siyuan Zhuang, Zhuohan Li, Danyang Zhuo, Stephanie Wang, Eric Liang, Robert Nishihara, Philipp Moritz, Ion Stoica.
35th Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM 2021).

HotOS'21 **In Reference to RPC: It's Time To Add Distributed Memory**.
Stephanie Wang, Benjamin Hindman, Ion Stoica.
18th Workshop on Hot Topics in Operating Systems (HotOS 2021).

NSDI'21 **Ownership: A distributed futures system for fine-grained tasks**.
Stephanie Wang, Eric Liang, Edward Oakes, Benjamin Hindman, Frank Sifei Luan, Audrey Cheng, Ion Stoica.
18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2021).

EuroS&P'20 **Practical volume-based attacks on encrypted databases**.
Rishabh Poddar*, Stephanie Wang*, Jianan Lu, Raluca Ada Popa.
5th IEEE European Symposium on Security and Privacy (EuroS&P 2020).

SOSP'19 **Lineage stash: Fault tolerance off the critical path**.
Distinguished Artifact
Stephanie Wang, John Liagouris, Robert Nishihara, Philipp Moritz, Ujval Misra, Alexey Tumanov, Ion Stoica.
27th ACM Symposium on Operating Systems Principles (SOSP 2019).

OSDI'18 **Ray: A distributed framework for emerging AI applications**.
Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, Ion Stoica.
13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2018).

SOSP'17 **Verifying a high-performance crash-safe file system using a tree specification**.
Haogang Chen, Tej Chajed, Alex Konradi, Stephanie Wang, Atalay İleri, Adam Chlipala, M. Frans Kaashoek, Nickolai Zeldovich.
26th ACM Symposium on Operating Systems Principles (SOSP 2017).

HotOS'17 **Real-time machine learning: The missing pieces**.
Robert Nishihara, Philipp Moritz, Stephanie Wang, Alexey Tumanov, William Paul, Johann Schleier-Smith, Richard Liaw, Mehrdad Niknami, Michael I Jordan, Ion Stoica.
16th Workshop on Hot Topics in Operating Systems (HotOS 2017).

Industry Positions

2019–Present **Anyscale**, *Software Engineer*, Berkeley, CA.
- Series C startup valued at $1bil based on the open-source project Ray.
- Launched large-scale 100TB shuffle as part of the Ray Datasets library.
- Building a flexible and scalable distributed data loader for machine learning pipelines.

2016–Present **Ray project**, *Co-creator and Lead committer*, github.com/ray-project/ray.

Summer 2017 **Samsara**, *Software Engineering Intern*, San Francisco, CA.

Summer 2014 **Locu**, *Software Engineering Intern*, San Francisco, CA.

January 2014 **Diffeo**, *Software Engineering Extern*, Cambridge, MA.

Summer 2013 **KAYAK**, *Software Engineering Intern*, Concord, MA.

Teaching

2017, 2021 **University of California, Berkeley**, *Graduate Student Instructor*, Berkeley, CA.
- CS262A: Advanced Topics in Computer Systems (graduate), Fall 2021.
- CS162: Operating Systems and Systems Programming (undergraduate), Fall 2017.

2013-2016 **Massachusetts Institute of Technology**, *Teaching Assistant*, Cambridge, MA.
- 6.824: Distributed Systems (undergraduate + graduate), Spring 2016.
- 6.004: Computation Structures (undergraduate introduction to computer architecture), Fall 2015.
- 6.042: Mathematics for Computer Science (undergraduate introduction to discrete mathematics), Fall 2013.

Service and Outreach

June 2021 **ICML'21**, *Reviewer*, virtual due to COVID-19.
Workshop on Challenges in Deploying and Monitoring Machine Learning Systems.

February 2020, October 2016 **Bay Area Graduate Pathways to STEM (GPS)**, *Peer Advisor*, Berkeley, CA.
Conference for low-income, first-generation, and/or underrepresented minority students to provide guidance in preparing for graduate school.

Spring 2019, Spring 2017 **University of California, Berkeley**, *EECS Visit Days Area Coordinator*, Berkeley, CA.

- Co-coordinator for admitted graduate students in operating systems, networking, databases, and security.
- Organized research and social events with current faculty and students. Arranged one-on-one meetings and "office hours" with current students.
- Also acted as peer advisor and student host over multiple years.

Fall 2018 **Be A Scientist**, *Mentor*, King Middle School, Berkeley, CA.
Weekly mentorship program for 7th grade students at Berkeley public school. Worked with a group of three students to design and carry out their own science experiments.

Summer 2016, Summer 2015 **Middle East Entrepreneurs of Tomorrow (MEET)**, *Computer Science Instructor (2015), Head of Curriculum (2016)*, Jerusalem, Israel and Cambridge, MA, USA.
Non-profit program to bring together young (high school) Israeli and Palestinian leaders to create positive change through teamwork and education in technology and entrepreneurship.

Selected Talks

August 2022 **Large-scale data shuffle in Ray with Exoshuffle**, *Ray Summit 2022*, San Francisco, CA.

April 2022 **From Ray to Anyscale: Research, startups, and open-source software**, *UT Austin*, virtual due to COVID-19.

July 2021 **The need for serverless Python**, SciPy 2021, virtual due to COVID-19.

June 2021 **In Reference to RPC: It's Time To Add Distributed Memory**, 18th Workshop on Hot Topics in Operating Systems *(HotOS 2021)*, virtual due to COVID-19.

April 2021 **Ownership: A Distributed Futures System for Fine-Grained Tasks**, *18th USENIX Symposium on Networked Systems Design and Implementation* (NSDI 2021), virtual due to COVID-19.

March 2021 **Ownership: A Distributed Futures System for Fine-Grained Tasks**, *Duke University*, Systems and Networking Seminar, virtual due to COVID-19.

February 2021 **Fault Tolerance in Ray: A General-Purpose Distributed System**, *UC Berkeley*, Berkeley EECS Annual Research Symposium (BEARS), virtual due to COVID-19.

October 2020 **Ray Internals: A Peek at 'ray.get'**, *Ray Summit 2020*, virtual due to COVID-19.

November 2019 **Distributed Memory for Futures with Ownership**, *IBM Research Almaden*, IBM Research Student Workshop on Systems and Cloud, San Jose, CA.

October 2019 **Lineage Stash: Fault Tolerance Off the Critical Path**, *27th ACM Symposium on Operating Systems Principles* (SOSP 2019), Huntsville, Ontario, Canada.

September 2018 **Ray: A Distributed System for Emerging AI Applications**, *Strange Loop*, with Robert Nishihara, St. Louis, MO.

Additional Media

whitepaper **Ray v2 Architecture**, *October 2022*.
tinyurl.com/ray-v2-whitepaper

blog **Analyzing memory management and performance in Dask-on-Ray**, *June 2021*.
https://medium.com/p/930a2236b70d

blog **Executing a distributed shuffle without a MapReduce system**, *March 2021*.
https://medium.com/p/d5856379426c

blog **Data Processing Support in Ray**, *February 2021*.
https://medium.com/p/ae8da34dce7e
With the Ray authors.

Awards

2021 Participant of EECS Rising Stars: Selective workshop for graduate students of marginalized or underrepresented genders in preparation for faculty careers

2020 Greylock X Fellow: Selective program for students interested in entrepreneurship

2019 Distinguished Artifact Award at SOSP'19

2016–2018 UC Berkeley Chancellor's Fellowship

2017 NSF Graduate Research Fellowship Program: Honorable Mention

2016 UC Berkeley EECS Excellence Award

2015 Phi Beta Kappa, MIT