

# Curriculum Vitae for Weijia Song

Senior Research Associate  
Department of Computer Science  
Cornell University

Phone: (917)-325-6960  
Email: [wsong@cornell.edu](mailto:wsong@cornell.edu)

---

## Research Interests

My research focuses on high-performance application frameworks for low-delay cloud AI/ML applications. I'm fascinated with pushing system performance to hardware limitations. I'm currently building a cloud system with optimized RDMA data paths called Cascade (<http://github.com/derecho-project/cascade>).

## Education

- Ph.D. in Computer Science, January 2014, Peking University  
Advisor: Zhen Xiao
- M.S. in Computer Science, April 2004, Beijing Institute of Technology
- B.S. in Computer Science, July 2001, Beijing Institute of Technology

## Employment

- November 2022 – Present: Senior Research Associate, Department of Computer Science, Cornell University. Research on distributed system and cloud computing.
- March 2017 – November 2022: Research Associate, Department of Computer Science, Cornell University. Research on distributed system and cloud computing.
- May 2014 – March 2017: Postdoctoral Researcher, Department of Computer Science, Cornell University. Research on distributed system and cloud computing.
- August 2006 – April 2014: Senior research and development staff member, Computer Center, Peking University. Working on network related research and software development.
- April 2004 – July 2006: Research and development staff member, Computer Center, Peking University. Working on network related research and software development.
- October 2000 – June 2001: Software engineer intern, Global Software Group, Motorola Co. Ltd. Working on software development for embedded systems.

## Publications

### Journal and Conference Papers (Citations: 1762, h-index: 10)

- **Song, Weijia**, et al. "Cascade: A Platform for Delay-Sensitive Edge Intelligence." *arXiv preprint arXiv:2311.17329* (2023).
- Garrett, T., **Song, W.**, Vitenberg, R., & Birman, K. (2023). Low-Latency ML Inference by Grouping Correlated Data Objects and Computation. *arXiv preprint arXiv:2312.11488*.
- Edward Tremel, **Weijia Song**, Ken Birman, Trust Without Waste: High-Throughput Tamper-Proof Distributed Logs with CascadeChain, *submitted to ICDCS 2024*
- **Song, Weijia**, et al. "Cascade: An Edge Computing Platform for Real-time Machine

Intelligence." Proceedings of the 2022 Workshop on Advanced tools, programming languages, and Platforms for Implementing and Evaluating algorithms for Distributed systems. 2022.

- Pengze Li, Lichen Pan, Xinzhe Yang, **Weijia Song**, Zhen Xiao, Ken Birman, Stabilizer: Geo-Replication with User-defined Consistency, *ICDCS 2022*
- Lorenzo Rosa, **Weijia Song**, Luca Foschini, Antonio Corradi, Ken Birman, DerechoDDS: Strongly Consistent Data Distribution for Mission-Critical Applications, *MILCOM 2021*
- **Song, Weijia**, Christina Delimitrou, Zhiming Shen, Robbert Van Renesse, Hakim Weatherspoon, Lotfi Benmohamed, Frederic De Vault, and Charif Mahmoudi. "CacheInspector: Reverse Engineering Cache Resources in Public Clouds." *ACM Transactions on Architecture and Code Optimization (TACO)* 18, no. 3 (2021): 1-25.
- Tremel, Edward, Sagar Jha, **Weijia Song**, David Chu, and Ken Birman. "Reliable, Efficient Recovery for Complex Services with Replicated Subsystems." In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pp. 172-183. IEEE, 2020.
- **Weijia Song**, Matthew Milano, Sagar Jha, Edward Tremel, Xinzhe Yang, and Ken Birman; Derecho's Extensible, Intelligent Object Store. AI Systems, ACM Symposium on Operating System Principles (SOSP) 2019
- Jha, S., Behrens, J., Gkountouvas, T., Milano, M., **Song, W.**, Tremel, E., ... & Birman, K. P. (2019). Derecho: Fast State Machine Replication for Cloud Services. *ACM Transactions on Computer Systems (TOCS)*, 36(2), 4.
- ANDERSON D, SAHASRABUDDHE S, BOSE A, HAUSER C, LITVINOV E, LUO X, ZHANG F, GKOUNTOUVAS T, **SONG W**, LIAO Y, BIRMAN K. Cloud-based Data Exchange Infrastructure for Wide Area Monitoring of Bulk Electric Power Grids. International Council on Large Electric Systems(CIGRE) 2018
- Robbert van Renesse, Hakim Weatherspoon, Zhiming Shen, and **Weijia Song**. The Supercloud: Applying Internet Design Principles to Interconnecting Clouds. *IEEE Internet Computing*. Jan./Feb. 2018.
- S. Ren, L. Tan, C. Li, Z. Xiao and **W. Song**, "Leveraging Hardware-Assisted Virtualization for Deterministic Replay on Commodity Multi-Core Processors," in *IEEE Transactions on Computers*, vol. 67, no. 1, pp. 45-58, Jan. 1 2018.  
doi: 10.1109/TC.2017.2727492
- Zhiming Shen, Qin Jia, Gur-Eyal Sela, **Weijia Song**, Hakim Weatherspoon, and Robbert Van Renesse. 2017. Supercloud: A Library Cloud for Exploiting Cloud Diversity. *ACM Trans. Comput. Syst.* 35, 2, Article 6 (October 2017), 33 pages. DOI: <https://doi.org/10.1145/3132038>
- **Weijia Song**, Theodoros Gkountouvas, Ken Birman, Qi Chen, and Zhen Xiao. FFFS: The Freeze-Frame File System. ACM Symposium on Cloud Computing (SoCC '16).
- Zhiming Shen, Qin Jia, Gur-Eyal Sela, Ben Rainero, **Weijia Song**, Robbert van Renesse, and Hakim Weatherspoon. Follow the Sun through the Clouds: Application Migration for Geographically Shifting Workloads. ACM Symposium on Cloud Computing (SoCC '16).
- Shiru Ren, Le Tan, Chunqi Li, Zhen Xiao, and **Weijia Song**. "Samsara: Efficient Deterministic Replay in Multiprocessor Environments with Hardware Virtualization Extensions." In 2016 USENIX Annual Technical Conference (USENIX ATC 16). 2016.
- Qin Jia, Zhiming Shen, **Weijia Song**, Robbert van Renesse, and Hakim Weatherspoon. "Smart spot instances for the supercloud." In Proceedings of the 3rd Workshop on CrossCloud

Infrastructures & Platforms, p. 5. ACM, 2016.

- Ian Kash, Qin Jia, Zhiming Shen, **Weijia Song**, Robbert van Renesse, and Hakim Weatherspoon. "Economics of a supercloud." In Proceedings of the 3rd Workshop on CrossCloud Infrastructures & Platforms, p. 7. ACM, 2016.
- Qin Jia, Zhiming Shen, **Weijia Song**, Robbert van Renesse, and Hakim Weatherspoon. "Supercloud: Opportunities and challenges." ACM SIGOPS Operating Systems Review 49, no. 1 (2015): 137-141.
- **Weijia Song**, Zhen Xiao, Qi Chen, and Haipeng Luo. Adaptive Resource Provisioning for the Cloud Using Online Bin Packing. IEEE Transactions on Computers (TC). Nov. 2014.
- Zhen Xiao, **Weijia Song**, and Qi Chen. Dynamic Resource Allocation using Virtual Machines for Cloud Computing Environment IEEE Transactions on Parallel and Distributed Systems (TPDS), June 2013.
- Qingnan Lai, Hao Ma, **Weijia Song**, Tingting Li, Guangxue Jiang, and Bei Zhang. Communication between Microblog with BBS and Analysis of User's Behavior in Social Network. Journal on Communications, 2013(Z1)
- **Weijia Song**, Hao Ma, Zhen Xiao, Xiaojun Zhang, and Bei Zhang. A Survey of Resource Scheduling for Virtualized Data Center. Journal of Guangxi University (Natural Science Edition), 2011(S1)
- **Weijia Song**, Changling Zhou, Jian Cui, Qun Shang, and Bei Zhang. Dynamic Dual-Binding: A Solution to ARP Cache Poisoning. Periodical of Ocean University of China, 2011(S1)
- Jia Yang, Hao Ma, **Weijia Song**, Jian Cui, and Changling Zhou. Crawling the eDonkey network. In Proceedings of 5<sup>th</sup> International Conference on Grid and Cooperative Computing Workshops, 2006.
- Jia Yang, **Weijia Song**, Bei Zhang, Xiaonan Li, and Xiaojun Zhang. Authentication and accounting based on IP control gateway and 802.1x in a campus network. Journal on Communications, 2006(Z1)
- Zhixue Lu, Hao Ma, **Weijia Song**, Jia Yang, and Bei Zhang. Study and Implementation on Visualization of P2P Systems. Journal on Communications, 2006(Z1)
- **Weijia Song**, Qun Shang, and Ping Chen. Designing and Implementing a Campus Network Resources Management System. Journal on Communications, 2005(Z1)
- Lifan Zhang, Xiaohua Wang, Jingsong Hu, **Weijia Song**, and Bin Long. Spatial Indices Based on Grid Partition. Journal of Beijing Institute of Technology, 2004(2).
- Jingsong Hu, Lifan Zhang, Xiaohua Wang, **Weijia Song**, and Bin Long. New Algorithm for Simple Polygon and Point. Computer Engineering, 2004(20).
- **Weijia Song**, Lifan Zhang, Xiaohua Wang, Wei Yan, Jingsong Hu, and Bin Long. A Road-topology Generation Algorithm Based on Skeletonization. Computer and Communications, 2004(3).

### **Doctoral Dissertation**

- **Weijia Song**, Resource Scheduling for Virtualized Data Center Environment (Doctoral Dissertation), Retrieved from Peking University Library(2013).

### **Book Chapters**

- **Weijia Song** and Zhen Xiao. An Infrastructure-as-a-Service Cloud: On-Demand Resource Provisioning. Invited book chapter in Principles, Methodologies, and Service-Oriented Approaches for Cloud Computing (pp. 301-323). Hershey, PA: Business Science Reference.

## Patents

- Zhen Xiao, **Weijia Song**, Qi Chen, Haipeng Luo, and Shuwei Chen. An Approach to Virtual Machine Resource Scheduling. Chinese Patent No. ZL 2009 1 0238364.2, Date of Grant: 10/10/2012
- **Weijia Song**, Hao Ma, Jianyu Zhang, Yuan Zhang, Jia Yang, Bei Zhang, and Yuan Zhou. An Approach to Crawling BitTorrent Seed Files. Chinese Patent No. ZL 2010 1 0147527.9, Date of Grant: 12/21/2011
- Zhen Xiao, Qun Huang, and **Weijia Song**. An Adaptive Load Predictor for Cloud Computing. Chinese Patent No. ZL 2011 1 0084332.9, Date of Grant: 11/21/2012
- Jia Yang, Hao Ma, **Weijia Song**, Bei Zhang, and Yuan Zhou. An Approach to Crawling the eDonkey Network. Chinese Patent No. ZL 2010 1 0263210.1, Date of Grant: 5/2/2012
- **Weijia Song**, Hao Ma, Jia Yang, Jian Cui, Ping Chen, and Bei Zhang. A Secure Group Communication Method. Chinese Patent No. ZL 2010 1 0531051.9, Date of Grant: 5/29/2013

## Conference and Colloquium Presentations

- “The Freeze-Frame File System” and “Follow the Sun through the Clouds: Application Migration for Geographically Shifting Workloads”, IEEE International Cloud Resiliency Workshop, Shenzhen, China, October 21<sup>st</sup>, 2016
- “The Freeze-Frame File System”, invited presentation at Futurewei Technologies, Santa Clara CA, October 14<sup>th</sup>, 2016
- “The Freeze-Frame File System”, SoCC2016, Santa Clara CA, October 6<sup>th</sup>, 2016
- “SuperCloud Demo”, NIST Cloud Computing Workshop, September, 2016
- “Follow the Sun through the Clouds: Application Migration for Geographically Shifting Workloads”, IEEE International Cloud Resiliency Workshop, Shenzhen, China, October 21<sup>st</sup>, 2016
- “Dynamic Resource Allocation using Virtual Machine for Cloud Computing Environment”, ChinaSys, Wuhan, China, October 2012
- “Building an IaaS System in Peking University”, Joint CHAIN/EPIKH School for Application Porting to Science Gateways, Beijing, China, April 2012
- “Building an IaaS System in Peking University”, The 5<sup>th</sup> AEARU Workshop on Network Education, Nanjing University, Nanjing, China, December 2011
- “Application Resources Integration to CARS Federation”, The 26<sup>th</sup> Asia-Pacific Advanced Network (APAN) Meeting, Queenstown, New Zealand, August 2008
- “IPv6 Home Network”, Technical Reports of China-Japan IPv6 Joint Project, Hangzhou and Wuhan, China, December 2005
- “IPv6 Based Home Gateway”, The 12<sup>th</sup> CERNET Annual Conference, Dalian, China, November 2005

## Research and Engineering Experiences

Department of Computer Science, Cornell University

**Senior Research Associate (Nov 2022 -now)**

**Research Associate (March 2017 – Oct 2022)**

**Postdoc Researcher (May 2014 – March 2017)**

Selected achievements:

- *June 2019 – now* As the leader of the Cascade project, I developed a novel C++17 cloud

application framework powered by optimized RDMA data paths. It provides a K/V API for data manipulation in distributed memory and persistent storage. It features high-throughput and low latency from zero-copy RDMA and NVMe storage layer; Its timestamp-indexed versioning capability allows reproducing system states anytime in the past. This open-source application is now available on GitHub.

- **November 2016 – now** As the lead developer of the Cross Data Site Replication for Cloud Storage Service project, I explored the cutting-edge consensus protocols to guarantee the strong consistency among data replicas across multiple datacenters. I came up with a two-tiered replication scheme enabling high availability as well as high performance. Based on this technology, we are developing a multi-tenant distributed data storage.
- **March 2017 – now** As one of the core developer of Derecho project - an RDMA based consensus protocol magnitude faster than existing solutions, I'm leading on two research subtopics. One goal is to improve the data replication performance/reliability in geographically distributed scenarios. The other direction is to enhance cloud-platform resource management capabilities (e.g., CPU and cache allocation), which would allow applications to balance the task objectives by changing behavior intelligently. In this ongoing project, I designed and implemented the high-performance persistent layer and RDMA communication layer in Derecho.
- **September 2016 – now** As the leader of *CacheInsight* project, I designed and implemented an efficient and software only cache probing tool that enables the Cloud users to determine the amount of CPU cache allocation, which usually change by time and generally not visible to the user. This tool enables many new possibilities. Time-critical applications could adapt to the dynamic allocation without violating deadline SLA. Cloud users have more insight into the services from various cloud providers and determine the optimal solution for their requirement.
- **May 2014 – 2017** As the lead developer of the Freeze-Frame File System (FFFS) project, I designed and implemented a distributed file system supporting temporal analysis by light-weighted precise and consistent snapshot mechanism, which is appealing to critical Cloud-based real-time applications like intelligent transportation system, smart meter. It currently achieves ten's of milliseconds snapshot precision and 450MBps read/write throughput with 10g Ethernet. We build FFFS with technologies including Remote DMA (RDMA), Hybrid Logical Clock, Network tuning, and log structured storage.
- **May 2014 – 2017** As the developer of the SuperCloud project, I designed and implemented the dynamic resource scheduling subsystem, which dynamically allocates cloud resources to fit demands. I also developed the Super-ZooKeeper application which can geographically follow the workload to achieve both high performance and low cost. (Supercloud is a Library Cloud that provides a user-level VM migration primitive, and can glue together one or more availability zones within a cloud provider or across different cloud providers.)
- **May 2014 – now** I am responsible for maintenance of Fractus, a 38-node private cloud service for computer system research at the department of computer science. This work included managing the Eucalyptus based virtualized cluster, adding new RDMA card and routers, upgrading drivers/operating system, and troubleshooting.

Computer Center, Peking University, Beijing, China

**Research and Development Staff Member (August 2006 – April 2014)**

Selected achievements:

- **August 2006 - April 2014** Developed the IPGW Client software (available at <http://its.pku.edu.cn/ipgwclient>) in the year 2006, which provides Internet access for the students, faculty, and staffs at Peking University. Continuously maintained it by adding features including ARP spoofing protection, multi-language support, and diagnostic tools. Ported it to popular operating systems such as Windows, Linux, Mac OS, iOS, Android, and Chrome browser. The peak number of online users is over ten thousand.
- **October 2009 - March 2011** Designed and deployed a secure communication overlay network based on Host Identity Protocol(HIP) over the open Internet. This project is supported by the Chinese National Science & Technology Pillar Program.
- **March 2010 - December 2010** Designed and deployed the IaaS Cloud System for the Institute of Network Computing and Information System, Peking University, to support the demand on virtual machines from over 200 students and research staffs in the institute. This system is built using OpenStack over 56 high-performance blade servers and 60 terabytes NAS storage.
- **December 2004 - July 2007** Designed an IPv6-enabled home gateway hardware based on Intel ixp425 network processor. Ported the Linux kernel 2.6.10 and the Java virtual machine to the platform. Developed software that manages the home appliances using the UPnP and the ITophome protocol (the Chinese home network standard). Established an open source project uriparser (available at <http://sourceforge.net/projects/uriparser/>) that implementing RFC 3986. Users of uriparser include Google and VMware. Added IPv6 support to the libupnp open source project (available at <http://sourceforge.net/p/pupnp/patches/1/>). This project is sponsored by China-Japan IPv6 Joint Project and China Next Generation Internet Project (CNGI).

Computer Center, Peking University, Beijing, China

#### **Assistant Research and Development Staff Member (April 2004 – July 2006)**

Selected achievements:

- **April 2004 - December 2004** Developed a network resource management system to support operation of the Peking University campus network. The system has been managed more than 6,000 devices and 170,000 IP addresses so far. It has been supporting the Network Operation Center at Peking University for more than thirteen years.

Department of Computer Science, Beijing Institute of Technology

#### **Graduate Student (September 2001 – July 2004)**

Selected achievements:

- **February 2002 - October 2003** Developed a vehicle GPS navigation system, an embedded system running above PC104 hardware. As the project leader of a four-member team, I was responsible for the architecture design, road topology extraction, and dead reckoning algorithm.

Global Software Group, Motorola, Beijing, China

#### **Software Engineer-Intern (October 2000 – June 2001)**

Selected achievements:

- **October 2000 - June 2001** Ported a ROM-based debugger monitor called MetroTRK to StarCore 140 DSP processor. It controls the code execution in the DSP processor using breakpoint instruction and talks to CodeWarrior debugger to realize debugging C/C++ and assembly code.

## Supervision of Students

Since 2014, I have supervised dozens of research projects for 35 Undergraduate and Master of Engineering (MEng) students. By doing this, I helped the students understand what a computer system research work is. Some outstanding students made significant contributions and became co-authors of our publications. The research projects include but are not limited to extending FFFS/Cascade/Derecho/CacheInspector with new features, evaluating/visualizing system performance, and experimenting with new technologies (SPDK, TensorFlow, Flink, etc.).

- Undergraduates '24: Noam Benson-Tilsen (nb498)
- Undergraduates '22: Aaron Jared Weiss (ajw242) and Archishman Sravankumar (as2647)
- Undergraduates '21: Aahil Suraj Awatramani (asa97), Zheng Wang (zw494), Liam Sullivan (lbs225), Benjamin Posnick (bmp53), Thompson Liu (ml2579), and Zicen Pan (zp37)
- Undergraduates '20: Alicia Yang (yy354), Samantha Deng (yd246), and Jeremy Ryan Glassman(jrg348)
- Undergraduates '19: Xinzhe Yang (xy269), Alex Katz (aik49), and Aayush Beri (ab2466)
- Undergraduates '18: Robert Barrett (rpb83) and Jordan Andrew Greissman (jag538)
- Undergraduates '17: Hannah Atmer (ha237), Nerla Jean-Louis (nj99), and Sacheth Sharath Hegde (ssh88)
- Undergraduates '16: Ben Rainero de Haan (br266) and William Jerome Schultz (wjs255)
- MEng '19: Shubham Agrawal (sa2279)
- MEng '18: Lu Yang (ly77), Weiyou Dai (wd248), and Brett Joseph Clancy(bjc265)
- MEng '17: Daniel Liu (dl596), Maarij Tahir (mt785), and Gautam Ravi Ramaswamy(grr37)
- MEng '16: Samir Borle (smb435), Xiao (Alex) Tan (xt85), and Richard Quan (rq32)
- MEng '15: Vidhya Sankaran (vs444), Anil Kumar Puvvadi (asp97), Sneha Prasad (sh824), and Feiran Chen (fc254)

## Former Duties

- ACM International Conference on Architectural Support for Programming Languages and Operating Systems 2024 (ASPLOS 2024), Program Committee member (7 artifacts evaluated)
- Reviewer for Journal of Parallel and Distributed Computing (JPDC) on paper submission: Multi-versioning Hardware Transactional Memory for Fail-Operational Multithreaded Applications. May 31, 2021, Major revision on Nov 27, 2021
- ACM Symposium on Cloud Computing 2020 (SoCC '20), Program Committee member (10 papers reviewed)
- Reviewer for Journal of Parallel and Distributed Computing (JPDC) on paper submission: Efficient Dynamic Resource Allocation based on Credibility (EDRAC) in cloud computing. April 28, 2019
- Reviewer for Journal KSII Transactions on Internet and Information Systems (Impact factor: 0.452) on paper submission: Application Centric Virtual Machine Placements to Minimize Bandwidth Utilization in Datacenters. March, 2018
- The 48<sup>th</sup> IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2018), External Reviewer (5 papers)
- Reviewer for Journal IEEE/ACM Transactions on Networking (Impact Factor 3.376) on papers submission: Optimal Capacity Provisioning for Online Job Allocation with Hard Allocation

Ratio Requirement. November, 2017

- Reviewer for Journal Frontiers of Information Technology & Electronic Engineering (Impact Factor 0.622) on paper submission: A TOGAF-based Method for Migrating Web Applications to Clouds with MVC and Microservice Architectures. December, 2016
- Reviewer for Journal Real-Time Systems (The International Journal of Time-Critical Computing Systems) on paper submission (Impact Factor 2.2): An Effective Method for Load Balancing and Task Allocation Based on Labeling and Multi-Layered-Agents in Cloud Computing. November, 2016
- 2016 USENIX Annual Technical Conference, External Reviewer (8 papers)
- The 23<sup>rd</sup> International World Wide Web Conference, WWW '14 – PC member (4 papers reviewed)
- Reviewer for IEEE Transactions on Parallel and Distributed Systems (Impact Factor 4.181) on paper submission: DPRA: Dynamic Power-saving Resource Allocation for Cloud Data Center using Particle Swarm Optimization. June, 2015.

## **Awards and Honors**

- Excellent Staff of Peking University Award, 2012
- Excellent Student Award of Institute of Network Computing and Information Systems, Peking University, 2011
- Outstanding Graduate of Beijing Award, 2004
- Outstanding Thesis of BIT Award, 2004