

Xianan (Shannon) Zhang

Current Address: 9138-H Regents RD, La Jolla, CA 92037

Phone: (858) 922-7185, Email: xzhang@cs.ucsd.edu

Education

- **Ph.D.** in Computer Science and Engineering, University of California, San Diego, 2006 (*expected*)
- **M.S.** in Computer Science and Engineering, University of California, San Diego, March 2002
- **B.S.** in Computer Science, Peking University, Beijing, China, July 1999

Ph.D. Dissertation: “Fault-tolerant Grid Services”, supervised by Prof. Keith Marzullo
Research focused on design and implementation of fault-tolerant services on Grid environments, including study of tradeoff between performance and utilization of Grid facilities associated with building reliable grid services, integration of different techniques to make grid service states durable, and investigation of practical protocols for replicating nondeterministic grid services on asynchronous systems.

Honors and Awards

- Winner of the MinuteSort World record of disk-to-disk sorting, 2000~2002
- Excellent Student Scholarship, 1996~1997, Peking University
- *LegendTM* Scholarship, 1995~1996, Peking University

Research Experience

Research Assistant, University of California, San Diego Summer 2000 ~ Present

- Nondeterministic services reliability
Designed and implemented practical protocols to replicate nondeterministic services – such as nondeterministic scheduling service and data replica selection service – on asynchronous systems.
- Performance study of Classic Paxos and Fast Paxos
Identified the factors that affect the practical performance of Classic Paxos and Fast Paxos, which enable efficient implementations of fault-tolerant state machines in asynchronous environments; demonstrated how these factors contribute to the performance penalty of each variant of Paxos using both probabilistic analysis and experiments with both a local cluster and PlanetLab.
- Customizable durability for Service Oriented Architectures (SOAs)
Designed an architecture that allows different durability mechanisms (e.g., in-memory replication or database storage) to be implemented as separate modules and then applied as needed to different variables and data structures constituting the service state, in order to optimize tradeoffs relative to runtime performance overhead, resource cost, and level of assurance provided.
- Fault-tolerance analysis of the GriPhyN system – a scientific DataGrid environment
Identified the fault-tolerance requirements of the GriPhyN system, which provides Grid techniques for scientific and engineering projects that must collect and analyze distributed, petabyte-scale datasets; designed fault-tolerant DAGMan which reliably schedules tasks submitted to the GriPhyN system.
- Highly-available services on OGSi and Globus

- Investigated the complexity and efficiency of building a reliable and highly available Grid service using the primary-backup replication on the Open Grid Services Infrastructure (OGSI) and its Globus toolkit implementations.
- Group membership services on wide-area systems
Analyzed the performance overhead of Group membership services on wide-area networks.
- MicroGrid API
Designed and implemented the client API of MicroGrid – online simulation tools for Grid applications.
- Disk-to-Disk sorting on High Performance Virtual Machines (HPVM)
Designed and implemented high performance algorithms for disk-to-disk sorting on clusters and held the minute sorting World record during 2000-2002.

Summer Intern, AT&T Labs – Research, Florham Park, NJ Summer 2003

- Grid fault-tolerance
Designed fault-tolerant architectures for Grid services based on the investigation on new Grid standards and existing fault tolerance techniques including J2EE and WS-Transaction standards.

Research Assistant, Peking University, Beijing, China Spring 1999

- “WebGather” machine learning system
Designed and implemented a system capable of automatically learning new Chinese words and phrases based on search logs for “WebGather”: A Chinese and English Web Search Engine.
- “WebGather” system tools
Built tools to collect and analyze system logs for “WebGather”.

Teaching Experience

Lecturer, University of California, San Diego Summer, 2004

- Taught undergrad course *Principles of Computer Operating Systems* (CSE120)

Teaching Assistant, University of California, San Diego Spring, 2004

- Led discussion sessions of undergrad course *Concurrency* (CSE128)

Publications

- **Xianan Zhang**, Flavio Junqueira, Matti A. Hiltunen, Keith Marzullo, Richard D. Schlichting, “Replicating Nondeterministic Services on Grid Systems”, submitted to the 15th IEEE International Symposium on High-Performance Distributed Computing (HPDC-15), 2006.
- **Xianan Zhang**, Flavio Junqueira, Keith Marzullo, “Fast Paxos vs. Classic Paxos: Caveat Emptor”, submitted to the 2006 International Conference on Dependable Systems and Networks (DSN-2006).
- **Xianan Zhang**, Matti A. Hiltunen, Keith Marzullo, and Richard D. Schlichting, “Customizable Service State Durability for Service Oriented Architecture”, submitted to the 3rd International Service Availability Symposium (ISAS 2006).
- **Xianan Zhang**, Dmitrii Zagorodnov, Matti A. Hiltunen, Keith Marzullo, and Richard D. Schlichting. “Fault-tolerant Grid Services Using Primary-Backup: Feasibility and Performance”, in Proceedings of the 2004 IEEE International Conference on Cluster Computing (Cluster 2004), September 2004.

- James Blythe, Richard Cavanaugh, Ewa Deelman, Ian Foster, Seung-Hye Jang, Cal Kesselman, Keith Marzullo, Reagan Moore, Valerie Taylor, and **Xianan Zhang**, “Types of Editors and Specifications”, GriPhyN technical report 2004-23, 2004.
- Kjetil Jacobsen, **Xianan Zhang**, and Keith Marzullo. “Group Membership and Wide-Area Master-Worker Computations”, in Proceedings of the 23rd International Conference on Distributed Computing Systems (ICDCS 2003), May 19-22, 2003, Providence, Rhode Island, USA.
- Hyojong Song, Xin Liu, Denis Jakobsen, Ranjita Bhagwan, **Xianan Zhang**, Kenjiro Taura, and Andrew Chien. “The MicroGrid: a Scientific Tool for Modeling Computational Grids”, Super Computing 2000 (SC 2000), Nov. 4-10, 2000, Dallas, TX.
- **Xianan Zhang**, Luis Rivera, and Andrew Chien, “HPVM MinuteSort”, Tech Report, UCSD, 2000.

Reference

Dr. Keith Marzullo
 Professor and Chair of Computer Science
 and Engineering Department
 University of California, San Diego
 Email: marzullo@cs.ucsd.edu
 Phone: (858) 534-3729

Dr. Richard Schlichting
 ACM/IEEE Fellow,
 Director of Software Systems Research
 AT&T Labs – Research
 Email: rick@research.att.com
 Phone: (973) 360-8234

Dr. Matti A. Hiltunen
 Research staff of Dependable Distributed
 Computing Department
 AT&T Labs – Research
 Email: hiltunen@research.att.com
 Phone: (973) 360-5504

Dr. Geoffrey M. Voelker
 Professor of Computer Science and
 Engineering Department
 University of California, San Diego
 Email: voelker@cs.ucsd.edu
 Phone: (858) 822-3323