

# Curriculum Vitae for Zhen Xiao

Professor  
Department of Computer Science  
Peking University  
Beijing, China, 100871

Phone: 010-62765815 ext. 8001 (China)  
Fax: 010-62765813 (China)  
E-mail: xiaozhen@pku.edu.cn  
Homepage: <http://zhenxiao.com>

---

## Research Interests

My research interests span several areas of distributed systems and machine learning, including cloud computing, big data, virtualization, fault-tolerant computing, security and dependability.

## Education

- Ph.D. in Computer Science, January 2001, Cornell University.  
Dissertation: Efficient Error Recovery for Reliable Multicast  
Advisor: Kenneth P. Birman  
GPA: 4.12/4.00 (A+ = 4.3, A = 4.0, A- = 3.7)
- M.S. in Computer Science, May 1999, Cornell University.
- B.S. in Computer Science, July 1996, Peking University, P.R. China.

## Awards and Honors

- **Teaching Excellence Award:** Peking University, September 2010.
- **Best Paper Award:** the 17th International workshop on Network and Operating Systems Support for Digital Audio & Video (NOSSDAV'07), June 2007.
- **Best Student Paper Award:** the 13th International World Wide Web Conference (WWW'04), May 2004.
- **Research Excellence Award:** AT&T Labs – Research, November 2002.

## Employment

- February 2008 – now: Professor, Department of Computer Science, Peking University. Working on cloud computing and big data.
- March 2006 – February 2008: Research Staff Member. IBM T.J. Watson Research Center, Hawthorne, NY.  
Worked on network system software.
- January 2001 – March 2006: Senior Technical Staff Member. AT&T Labs – Research, Florham Park, NJ.  
Worked on distributed systems research.
- August 1996 – December 2000: Research/Teaching assistant. Cornell University, Ithaca, NY.  
Worked on fault-tolerant computing.

## Research Grants

I served as the Principal Investigator (PI) for several grants with a total amount of about four million RMB. Besides, I have also participated as the co-PI in two 973 projects with a total amount of more than nine million RMB.

- “ Research on a Distributed Machine Learning Framework ”, National Natural Science Foundation of China, RMB 752,000
- “Research on Pattern Discovery and Impact Analysis of Network Big Data”, National Grand Fundamental Research 973 Program of China, RMB 6.3 million

- “Research on an Information Platform for Talent Management and Evaluation”, Beijing Municipal Commission of Education, RMB 150,000
- “Intelligent Storage Architecture and Technology for Big Data Analytics”, National High Technology Research and Development Program (“863” Program) of China, RMB 1.1 million
- “A Distributed Resource Sharing Platform on the Next Generation Internet”, Beijing Municipal Commission of Education, RMB 161,300
- “Efficient Resource Provisioning for Cloud Computing Services”, Okawa Foundation Research Grant, JPY one million
- “On-line deterministic replay of virtual machines (VMs) in multi-core environment”, National Natural Science Foundation of China, RMB 550,000
- “Large Scale Video-on-demand Application for the Chinese Next Generation Internet”, CNGI, RMB 680,000
- “Resource Provisioning Strategies for Cloud Computing”, Ministry of Education (MoE)-Intel Joint Research Fund, US \$40,000
- “Research on System Virtualization”, National Grand Fundamental Research 973 Program of China, RMB 3 million
- “Resource Scheduling in Cloud Computing Environment”, National Innovation Plan, RMB 10,000
- “Research on the Fundamental Theory and Application of Big Data Computing”, National Grand Fundamental Research 973 Program of China, RMB 6.3 million

## Professional Activities

- **Program Committees, Vice/Track Chairs**

1. The International Conference on Distributed Computing Systems (ICDCS): 2004, 2007(Vice Chair), 2009, 2010, 2011, 2012, and 2017.
  2. The 2016 USENIX Annual Technical Conference
  3. The International World Wide Web Conference (WWW): 2005, 2008, 2009, 2011, 2013, and 2014 (Track Chair)
  4. The International Conference on Dependable Systems and Networks (DSN): 2010 and 2011
  5. The International Parallel and Distributed Processing Symposium (IPDPS 2014)
  6. The International Conference on Cloud Computing Technology and Science (CloudCom 2013)
  7. The International Wireless Communications and Mobile Computing Conference (IWCMC 2006)
  8. The International Conference on Parallel Processing (ICPP 2005)
  9. IEEE Infocom 2005
  10. The International Conference on Parallel and Distributed Systems (ICPADS 2004)
  11. The International Workshop on Web Content Caching and Distribution (WCW 2003)
- ACM and IEEE Senior Member
  - Program co-chair for Internet Protocol TV workshop, jointly held with the 15th International World Wide Web Conference (WWW’06)
  - Served on various NSF panels in the U.S.
  - Panel Moderator for the 8th International Workshop on Web Content Caching and Distribution (WCW 2003): *Caching the Uncachable: An Uphill Battle Against Dynamic and Personalized Content*

## Citations

According to my profile in Google Scholar, the total number of citations for all my papers is 4695 (h-index 33) as of Feb 2018. Among those citations, 1898 are since year 2013.

## Selected Journal/Book Publications

The following publications can be downloaded from <http://zhenxiao.com/papers>

- Shiru Ren, Le Tan, Chunqi Li, Zhen Xiao, and Weijia Song. “Leveraging Hardware-Assisted Virtualization for Deterministic Replay on Commodity Multi-Core Processors”, *IEEE Transactions on Computers (TC)*, January 2018.
- Qi Chen, Jinyu Yao, Benchao Li, and Zhen Xiao. “PISCES: Optimizing Multi-job Application Execution in MapReduce”, *IEEE Transactions on Cloud Computing (TCC)*, August 2016.
- Qi Chen, Jinyu Yao, and Zhen Xiao. “LIBRA: Lightweight Data Skew Mitigation in MapReduce”, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, September 2015.
- Weijia Song, Zhen Xiao, Qi Chen, and Haipeng Luo. “Adaptive Resource Provisioning for the Cloud Using Online Bin Packing”, *IEEE Transactions on Computers (TC)*, November 2014.
- Zhen Xiao, Qi Chen, and Haipeng Luo. “Automatic Scaling of Internet Applications for Cloud Computing Services”, *IEEE Transactions on Computers (TC)*, May 2014.
- Qi Chen, Cheng Liu, and Zhen Xiao. “Improving MapReduce Performance Using Smart Speculative Execution Strategy”, *IEEE Transactions on Computers (TC)*, April 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.
- Jun Zhu, Zhefu Jiang, Zhen Xiao, and Xiaoming Li. “Optimizing the Performance of Virtual Machine Synchronization for Fault Tolerance”, *IEEE Transactions on Computers (TC)*, November 2011.
- Zhen Xiao. “Cloud computing and virtualization (in Chinese)”, *Communications of CCF (China Computer Federation)*, October 2009.
- Lei Guo, Songqing Chen, Zhen Xiao, Enhua Tan, Xiaoning Ding, and Xiaodong Zhang. “A Performance Study of BitTorrent-like Peer-to-Peer Systems”, *IEEE Journal on Selected Areas in Communication (JSAC), Special Issue on Peer-to-Peer Communications and Applications*, Vol 25, No. 1, January 2007.
- Yennun Huang, Yih-Farn Chen, Rittwik Jana, Hongbo Jiang, Michael Rabinovich, Amy Reibman, Bin Wei, and Zhen Xiao. “Capacity Analysis of MediaGrid: a P2P IPTV Platform for Fiber to the Node (FTTN) Networks”, *IEEE Journal on Selected Areas in Communication (JSAC), Special Issue on Peer-to-Peer Communications and Applications*, Vol 25, No. 1, January 2007.
- Leeann Bent, Michael Rabinovich, Geoffrey M. Voelker, and Zhen Xiao. “Characterization of a Large Web Site Population with Implications for Content Delivery”, *World Wide Web Journal*, Springer Netherlands, Volume 9, Number 4, pages 505-536, December 2006.
- Christof Fetzer and Zhen Xiao. “Automatic Testing of Errors and Vulnerabilities in Software Components”, Invited book chapter in *Testing Commercial-off-the-shelf Components and Systems*, Springer-Verlag, January 2005.
- Kenneth P. Birman, Mark Hayden, Ozgur Ozkasap, Zhen Xiao, Mihai Budiu, and Yaron Minsky. “Bimodal Multicast”, *ACM Transactions on Computer Systems (TOCS)*, May 1999.

## Selected Conference/Workshop Publications

The following publications can be downloaded from <http://zhenxiao.com/papers>

- Chao Ma, Yuze Liao, Yuan Wang, and Zhen Xiao. “F2M: Scalable Field-Aware Factorization Machines”, *Proc. of the MLSys on NIPS 2016*, December 2016.

- Yang Xiao, Yuan Wang, Hangyu Mao, and Zhen Xiao. “Predicting Restaurant Consumption Level through Social Media Footprints”, *Proc. of the International Conference on Computational Linguistics (COLING 2016)*, December 2016.
- Yuan Wang, Yang Xiao, Chao Ma, and Zhen Xiao. “Improving Users’ Demographic Prediction via the Videos They Talk about”, *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP16)*, November 2016.
- Weijia Song, Theo Gkountouvas, Ken Birman, Qi Chen, and Zhen Xiao. “The Freeze-Frame File System”, *Proc. of the ACM Symposium on Cloud Computing (SoCC’16)*, October 2016.
- Shiru Ren, Le Tan, Chunqi Li, Zhen Xiao, and Weijia Song. “Samsara: Efficient Deterministic Replay in Multiprocessor Environments with Hardware Virtualization Extensions”, *Proc. of the USENIX Annual Technical Conference (ATC 2016)*, June 2016.
- Chao Ma, Yan Ni, and Zhen Xiao. “Brook: An Easy and Efficient Framework for Distributed Machine Learning”, *Proc. of the LearningSys Workshop on NIPS 2015*, December 2015.
- Shiru Ren, Chunqi Li, Le Tan, and Zhen Xiao. “Samsara: Efficient Deterministic Replay with Hardware Virtualization Extensions”, *Proc. of the ACM SIGOPS Asia-Pacific Workshop on Systems (APSys 2015)*, July 2015.
- Chien-Chin Huang, Qi Chen, Zhaoguo Wang, Russell Power, Jorge Ortiz, Jinyang Li, and Zhen Xiao. “Spartan: A Distributed Array Framework with Smart Tiling”, *Proc. of the USENIX Annual Technical Conference (ATC 2015)*, July 2015.
- Yang Xiao, Wayne Xin Zhao, Kun Wang, and Zhen Xiao. “Knowledge Sharing via Social Login: Exploiting Microblogging Service for Warming up Social Question Answering Web sites”, *Proc. of the International Conference on Computational Linguistics (COLING 2014)*, August 2014.
- Qi Chen, Yisheng Liao, Christopher Mitchell, Jinyang Li, and Zhen Xiao. “Building Scalable Multimedia Search Engine Using Infiniband”, *Proc. of the USENIX Workshop on Hot Topics in Cloud Computing (HotCloud)*, June 2014.
- Jun Zhu, Zhefu Jiang, and Zhen Xiao. “Twinkle: A Fast Resource Provisioning Mechanism for Internet Services”, *Proc. of IEEE Infocom*, April 2011.
- Jun Zhu, Wei Dong, Zhefu Jiang, Tiaogang Zhi, Zhen Xiao, and Xiaoming Li. “Improving the Performance of Hypervisor-Based Fault Tolerance”, *Proc. of the International Parallel and Distributed Processing Symposium (IPDPS 2010)*, April 2010.
- Zhi Yang, Yafei Dai, and Zhen Xiao. “Exploring the Cost-Availability Tradeoff in P2P Storage Systems”, *Proc. of the International Conference on Parallel Processing (ICPP 2009)*, September 2009.
- Lei Guo, Enhua Tan, Songqing Chen, Zhen Xiao, and Xiaodong Zhang. “The Stretched Exponential Distribution of Internet Media Access Patterns”, *Proc. of the ACM Symposium on Principles of Distributed Computing (PODC 2008)*, August 2008.
- Zhen Xiao and Fan Ye. “New Insights on Internet Streaming and IPTV”, *Proc. of the ACM International Conference on Image and Video Retrieval (CIVR 2008)*, July 2008.
- Guanhua Yan, Zhen Xiao, and Stephan Eidenbenz. “Catching Instant Messaging Worms with Change-Point Detection Techniques”, *Proc. of the USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET ’08)*, April 2008.
- Zhen Xiao, Lei Guo, and John Tracey. “Understanding Instant Messaging Traffic Characteristics”, *Proc. of the 27th International Conference on Distributed Computing Systems (ICDCS’07)*, June 2007.
- Yih-Farn Chen, Yennun Huang, Rittwik Jana, Hongbo Jiang, Michael Rabinovich, Bin Wei, and Zhen Xiao. “When is P2P technology beneficial to IPTV services?” *Proc. of the 17th International workshop on Network and Operating Systems Support for Digital Audio & Video (NOSSDAV’07)*, June 2007. **Best Paper Award**

- Lei Guo, Enhua Tan, Songqing Chen, Zhen Xiao, and Xiaodong Zhang. “Does Internet Media Traffic Really Follow Zipf-like Distribution?” *Proc. of ACM SIGMETRICS (Extended Abstract)*, June 2007.
- Lei Guo, Enhua Tan, Songqing Chen, Zhen Xiao, Oliver Spatscheck, and Xiaodong Zhang. “Delving into Internet Streaming Media Delivery: A Quality and Resource Utilization Perspective”, *Proc. of ACM SIGCOMM Internet Measurement Conference (IMC’06)*, October 2006.
- Lei Guo, Songqing Chen, Zhen Xiao, Enhua Tan, Xiaoning Ding, and Xiaodong Zhang. “Measurements, Analysis, and Modeling of BitTorrent-like Systems”, *Proc. of ACM SIGCOMM Internet Measurement Conference (IMC’05)*, October 2005.
- Lei Guo, Songqing Chen, Zhen Xiao, and Xiaodong Zhang. “DISC: Dynamic Interleaved Segment Caching for Interactive Streaming”, *Proc. of the 25th International Conference on Distributed Computing Systems (ICDCS’05)*, June 2005.
- Lei Guo, Songqing Chen, Zhen Xiao, and Xiaodong Zhang. “Analysis of Multimedia Workloads with Implications for Internet Streaming”, *Proc. of the 14th International World Wide Web Conference (WWW’05)*, May 2005.
- Leeann Bent, Michael Rabinovich, Geoffrey M. Voelker, and Zhen Xiao. “Characterization of a Large Web Site Population with Implications for Content Delivery”, *Proc. of the 13th International World Wide Web Conference (WWW’04)*, May 2004. **Best Student Paper Award**
- Michael Rabinovich, Zhen Xiao, and Amit Aggarwal. “Computing on the Edge: A Platform for Replicating Internet Applications”, *Proc. of the Eighth International Workshop on Web Content Caching and Distribution (WCW’03)*, September 2003.
- Christof Fetzer and Zhen Xiao. “HEALERS: A Toolkit for Enhancing the Robustness and Security of Existing Applications”, *Proc. of the International Conference on Dependable Systems and Networks (Practical Experience and Demonstrations) (DSN’03)*, June 2003.
- Jennifer Rexford, Jia Wang, Zhen Xiao, and Yin Zhang. “BGP Routing Stability of Popular Destinations”, *Proc. of ACM SIGCOMM Internet Measurement Workshop (Extended Abstract) (IMW’02)*, November 2002.
- Christof Fetzer and Zhen Xiao. “A Flexible Generator Architecture for Improving Software Dependability”, *Proc. of the 13th International Symposium on Software Reliability Engineering (ISSRE’02)*, November 2002.
- Pradnya Karbhari, Michael Rabinovich, Zhen Xiao, and Fred Douglis. “ACDN: A Content Delivery Network for Applications”, *Proc. of ACM SIGMOD (Project Demo)*, June 2002.
- Christof Fetzer and Zhen Xiao. “An Automated Approach to Increasing the Robustness of C Libraries”, *Proc. of the International Conference on Dependable Systems and Networks (DSN’02)*, June 2002.
- Zhen Xiao, Kenneth P. Birman, and Robbert van Renesse. “Optimizing Buffer Management for Reliable Multicast”, *Proc. of the International Conference on Dependable Systems and Networks (DSN’02)*, June 2002.
- Christof Fetzer and Zhen Xiao. “Detecting Heap Smashing Attacks Through Fault Containment Wrappers”, *Proc. of the 20th IEEE Symposium on Reliable Distributed Systems (SRDS’01)*, October 2001.
- Zhen Xiao and Kenneth P. Birman. “A Randomized Error Recovery Algorithm for Reliable Multicast”, *Proc. of IEEE Infocom*, April 2001.
- Zhen Xiao and Kenneth P. Birman. “Providing Efficient, Robust Error Recovery Through Randomization”, *Proc. of the International Workshop on Applied Reliable Group Communication (Jointly held with the 21st International Conference on Distributed Computing Systems (ICDCS’01))*, April 2001.

## Selected US Patents

1. Songqing Chen, Lei Guo, Zhen Xiao, and Xiaodong Zhang. “System and method for streaming media objects”. U.S. Patent Number 8,230,098, issued on July 24, 2012.

2. Songqing Chen, Xiaoning Ding, Lei Guo, Enhua Tang, Zhen Xiao, and Xiaodong Zhang. "Method and System for Scalable Content Storage and Delivery". U.S. Patent Number 8,224,968, issued on July 17, 2012.
3. Yennun Huang, Yih-Farn Chen, Rittwik Jana, Amy Reibman, Bin Wei, Zhen Xiao, and Michael Rabinovich. "System and Method for Peer to Peer Video Streaming". U.S. Patent Number 7,903,652, issued on March 8, 2011.
4. Claudia Canali, Alexandre Gerber, Stephen Fisher, Michael Rabinovich, Oliver Spatscheck, and Zhen Xiao. "Monitoring For Replica Placement and Request Distribution". U.S. Patent Number 7,941,556, issued on May 10, 2011.
5. Arup Acharya, Erich M. Nahum, John Tracey, Xiping Wang, Charles P. Wright, and Zhen Xiao. "Overload protection for SIP servers". U.S. Patent Number 7,522,581, issued on April 21, 2009.
6. Zhen Xiao and Christof Fetzer. "Method for Detecting Software Errors and Vulnerabilities". U.S. Patent Number 7,293,213, issued on November 6, 2007.
7. Christof Fetzer and Zhen Xiao. "Robust Software Library Wrapper Method and Apparatus". U.S. Patent Number 7,225,428, issued on May 29, 2007.
8. Christof Fetzer and Zhen Xiao. "Methods and Apparatus for Detecting Heap Smashing". U.S. Patent Number 6,832,302, issued on December 14, 2004.
9. Lei Guo, Erich M. Nahum, John Tracey, Dinesh Verma, Xiping Wang, Charles P. Wright, and Zhen Xiao. "Method and Apparatus for Processing Messages in Messaging System". U.S. Patent filed on August 2007.
10. Leeann Bent, Michael Rabinovich, Geoffrey Voelker, and Zhen Xiao. "Method and Apparatus for Prefetching Web Content". U.S. patent filed in August 10, 2005.
11. Songqing Chen, Lei Guo, Zhen Xiao, and Xiaodong Zhang. "Caching Method and System for Providing Interactive Streaming". U.S. patent filed on January 31, 2005.
12. Songqing Chen, Lei Guo, Zhen Xiao, and Xiaodong Zhang. "Method and System for Supplying Media over Communication Networks". U.S. patent filed on January 31, 2005.
13. Chuck Kalmanek, Michael Rabinovich, and Zhen Xiao. "Method for Determining Where Dynamic Contents Should Be Assembled". U.S. patent filed on December 29, 2003.
14. Fred Douglass, Michael Rabinovich, and Zhen Xiao. "Method for Client-Side Inclusion of Data Elements". U.S. patent filed on March 22, 2002.

## Conference and Colloquium Presentations

Slides for the following talks can be downloaded from <http://zhenxiao.com/talks>

1. "A System Perspective on Cloud Computing", Department of Computing and Information Systems seminar, University of Melbourne, Australia, Sept 2014.
2. "A Personal Perspective on Cloud Computing", Seminar, Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong, Dec 2012.
3. "Understanding the benefit and limitation of cloud computing: What can virtualization buy us?", System Research Seminar, Department of Computer Science, University of Madrid, Spain, Feb 2012.
4. "Twinkle: A Fast Resource Provisioning Mechanism for Internet Services". IEEE Infocom, Shanghai, China, April 2011.
5. "Gone with the cloud: Adaptive Resource Virtualization for Amazon EC2-like Environment", System Research Seminar, Department of Computer Science, Cornell University, November 2009.
6. "Understanding Instant Messaging Traffic Characteristics", the 27th International Conference on Distributed Computing Systems (ICDCS'07), Toronto, Canada, June 2007.

7. "Delving into Internet Streaming Media Delivery: A Quality and Resource Utilization Perspective", Invited Talk, AT&T Labs – Research, Feb 2007.
8. "Understanding Web content delivery performance: Why is the Web still slow and what can we do about it?", Invited Talk, Department of Computer Science, University of California at Berkeley, October 2005.
9. "Optimizing Internet Content Delivery Systems – A View from the Inside", Invited Talk, Department of Computer Science, University of Tsukuba, Japan, May 2005.
10. "Analysis of Multimedia Workloads with Implications for Internet Streaming", the 14th International World Wide Web Conference (WWW'05), Chiba, Japan, May 2005.
11. "An Automated Approach to Software Reliability and Security", Invited Talk, Department of Computer Science, University of California at Berkeley, June 2003.
12. "HEALERS: A Toolkit for Enhancing the Robustness and Security of Existing Applications", IEEE International Conference on Dependable Systems and Networks (DSN'03), San Francisco, California, June 2003.
13. "Moving Edge Side Includes to the Real Edge – the Clients", USENIX Symposium on Internet Technologies and Systems (USITS'03), Seattle, Washington, March 2003.
14. "Client-Side Includes: Faster Access to Dynamic Web Content", AT&T Innovation Forum, Middletown, New Jersey, October 2002.
15. "Optimizing Buffer Management for Reliable Multicast", IEEE International Conference on Dependable Systems and Networks (DSN'02), Bethesda, Maryland, June 2002.
16. "ACDN: A Content Delivery Network for Applications", ACM SIGMOD, Madison, Wisconsin, June 2002.
17. "A Randomized Error Recovery Algorithm for Reliable Multicast", IEEE Infocom, Anchorage, Alaska, April 2001.
18. "Providing Efficient, Robust Error Recovery Through Randomization", IEEE International Workshop on Applied Reliable Group Communication, Phoenix, Arizona, April 2001.