

Zhenhuan Gong

1401 Collegiate Cir, Apt 201 *E-mail:* steve.gongzhenhuan@gmail.com zgong@ncsu.edu
Raleigh, NC, 27606 *Phone:* (919) 946-2597 (C)
Homepage: <http://www4.ncsu.edu/~zgong/>

Education

North Carolina State University	Computer Science	Ph.D Student (GPA 4.0)	2008 – present
Nanjing University, Nanjing, China	Computer Science	B.S. (GPA 3.6)	2002 – 2006

Work Experience

Research/Teaching Assistant	North Carolina State University, NC, USA	2008 – present
Research Intern	IBM T.J. Watson Research Center, NY, USA	2010 Spring
Engineer	Tencent Technology, Shenzhen, China	2006 – 2008

Awards/Honors

- IBM PhD Fellowship Award, 2010
- IEEE CNSM 2010 Best Paper Award, 1 out of 176 submissions, 2010
- IEEE MASCOTS 2010 Best Paper Nominee, 6 out of 40 accepted papers, 2010
- Excellent Undergraduate Student of Nanjing University, 2006
- Award by China Mobile for Excellent Students in College Entrance Examination, 2002 (Rank Top 10 in the city of Nanjing and Top 200 in Jiangsu Province among 300,000 students)
- People's Scholarships, Second prize once, Third prize twice, 2002–2005
- Excellent Student of Dept. of Computer Science and Technology three times, 2002–2005
- Excellent Student in the Military Training of Nanjing University, 2002
- Excellent Volunteer of the 10th Chinese National Games held in Jiangsu Province, 2005

Research Interests

My general research interests are *High-Performance Computing* and *Cloud Computing*. Specifically, I am focusing on storage management for scientific data, and resource management for cloud systems. I am also interested in scientific data compression, stream processing systems and online anomaly diagnose for distributed systems.

Publications

- "Multi-level Layout Optimization for Efficient Spatio-temporal Queries on ISABELA-compressed Data",
Zhenhuan Gong, Sriram Lakshminarasimhan, John Jenkins, Hemanth Kolla, Stephane Ethier, Jackie Chen, Robert Ross, Scott Klasky, Nagiza F. Samatova, IEEE International Parallel and Distributed Processing Symposium (**IPDPS**), Shanghai, China, May, 2012. (To appear)
- "S-preconditioner for Multi-fold Data Reduction with Guaranteed User-controlled Accuracy",
Ye Jin, Sriram Lakshminarasimhan, Neil Shah, **Zhenhuan Gong**, and Nagiza F. Samatova, IEEE International Conference on Data Mining (**ICDM**), Vancouver, Canada, December, 2011.

- "ISABELA-QA: Query-driven Data Analytics over ISABELA-compressed Extreme-Scale Scientific Data",
Sriram Lakshminarasimhan, Jonathan Jenkins, Robert Latham, Robert Ross, Nagiza F. Samatova, Isha Arkatkar, **Zhenhuan Gong**, Hemanth Kolla, Jackie Chen, Seung-Hoe Ku, C.S. Chang, Stephane Ethier, Scott Klasky, IEEE/ACM Conference on Supercomputing (**SC**), Seattle, Washington, November, 2011.
- "PRESS: PRedictive ELastic ReSource Scaling for Cloud Systems",
Zhenhuan Gong, Xiaohui Gu, John Wilkes, IEEE International Conference on Network and Services Management (**CNSM**), Niagara Falls, Canada, October, 2010. (acceptance rate: $27/176 = 15\%$) *Best Paper Award*
- "PAC: Pattern-driven Application Consolidation for Efficient Cloud Computing",
Zhenhuan Gong, Xiaohui Gu, IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (**MASCOTS**), Miami Beach, Florida, August, 2010. (acceptance rate: 16%) *Best Paper Nominee*
- "Predictive Elastic Load Management for Cloud Computing Infrastructures",
Zhenhuan Gong, Xiaohui Gu, ACM Symposium on Operating Systems Principles (**SOSP**) Poster Session, Big Sky, MT USA, Oct, 2009.
- "SigLM: Signature-Driven Load Management for Cloud Computing Infrastructures",
Zhenhuan Gong, Prakash Ramaswamy, Xiaohui Gu, Xiaosong Ma, IEEE International Conference on Quality of Service (**IWQoS**), Charleston, South Carolina, July, 2009.
- "Self-Correlating Predictive Information Tracking for Large-Scale Production Systems",
Ying Zhao, Yongmin Tan, **Zhenhuan Gong**, Xiaohui Gu, Mike Wamboldt, IEEE International Conference on Autonomic Computing and Communications (**ICAC**), Barcelona, Spain, June, 2009. (acceptance rate: $15/96 = 15.6\%$)

Project Experience

- 2011 – 2012, NCSU, **Multi-level storage layout optimization for scientific datasets**
Designed and implemented a multi-level storage layout optimization infrastructure, and a parallel query engine based on the infrastructure for multi-dimensional double-precision scientific datasets, in supercomputing environments. The infrastructure provided hierarchical optimization for various data access patterns introduced by value-constrained and spatial-constrained data queries, multi-resolution data access and so on. It applied light-weight indexing and data compression to reduce I/O and storage overhead for large datasets.
- 2010 – 2011, IBM and NCSU, **Predictive anomaly prevention for Cloud systems**
Designed and implemented a dynamic performance anomaly diagnose and prevention system based on machine learning techniques and out-of-box resource perturbation. The system applied Bayesian network classifiers to learn and predict performance anomalies, and dynamically adjusted resource allocated to virtual machines to prevent the anomalies.
- 2009 – 2011, NCSU, **Elastic resource scaling for Cloud systems**
Designed and implemented the resource scaling system based on real-time out-of-box workload tracing, characterization and prediction for virtual machines. The project addressed real-time tracing and prediction, dynamic resource allocation and fast reaction to workload changes, to achieve best resource utilization for Cloud resources with guaranteed application performance.

- 2008 – 2010, NCSU, **Pattern-driven service consolidation for Cloud systems**
Designed and implemented the consolidation system based on out-of-box virtual machine monitoring, smart workload characterization and pattern matching algorithms. Built Xen-based virtualized Cloud infrastructure and monitoring systems, and the virtualized multi-tier RUBiS Web service benchmarking environment. Organized and ran large-scale experiments and collected performance results.
- 2006 – 2008, Tencent Technology, **Linux back-end development**
Worked in Tencent Technology as system developer for QQPet, a large interactive Web game community with over 1.5 million Peak Concurrent Users. The job includes design, development and maintenance of Web CGIs, database and back-end server applications in Linux environment. Very quick and proficient in resolving both technical and algorithm related problems under complex game logic.
- 2006, Nanjing University, **Machine learning-based XQuery optimization**
Took part in a research of performance and optimization of XQuery using the method of reinforcement learning with several graduate students in our department. Designed and implemented the main program using Java-based Apache Xerces Package and fatdog XML engine to demonstrate and compare the performance of different ways of structural joins.
- 2005, Nanjing University, **Java-based Instant Messenger**
Developed an ICQ-like Instant Messenger tool using JBuilder 9 and SQL Server 2000 under Windows. Designed and implemented the server side logic, network communication and DB architecture. The tool implements nearly all the basic functions of an instant messenger.

Computer Skills

- Comprehensive knowledge and rich experience in development in Linux environment.
- Good experience in Linux system administration and maintenance.
- Programming Languages: C, C++, Java, Python, Shell, Matlab, Perl, Javascript.
- System/Platform Experiences: Linux, MPI, TCP/IP, Web CGI, MySQL, HTML, AJAX, Xen, QEMU/KVM, Map/Reduce, IBM Infosphere, J2EE, Linux kernel modules and so on.

Miscellaneous

Citizen of China. Student F-1 visa.