

Xiaohui (Helen) Gu

890 Oval Drive
Department of Computer Science
North Carolina State University
Raleigh, NC 27695

E-mail: xiaohui.helen.gu@gmail.com
Phone: (919) 515-9676 (O), (914) 299-1153 (C)
Fax: (919) 515-7896
Homepage: <http://www.csc.ncsu.edu/faculty/gu/>

Research Interests

My general research interests are in the areas of systems, with a current focus on cloud computing, system reliability, and system integrity.

Education

University of Illinois at Urbana-Champaign, Urbana, IL	Computer Science	Ph.D.	2004
University of Illinois at Urbana-Champaign, Urbana, IL	Computer Science	M.S.	2001
Peking University, Beijing, China	Computer Science	B.S.	1999

Professional Experience

Assistant Professor	CS Dept, North Carolina State University	2008–present
Research Staff Member	IBM T.J. Watson Research Center, Hawthorne, NY	2004- 2007
Research Intern	IBM T.J. Watson Research Center, Hawthorne, NY	2003 Summer
Research Intern	IBM T.J. Watson Research Center, Hawthorne, NY	2002 Summer
Research Intern	Hewlett-Packard Research Labs, Palo Alto, CA	2001 Summer

Awards/Honors

- NSF CAREER Award, 2012.
- IBM Faculty Award, 2008, 2009, 2010, 2011.
- Google Research Award, 2009, 2011.
- IEEE CNSM 2010 Best Paper Award, 1 out of 176 submissions.
- IEEE MASCOTS 2010 best paper nominee, 6 out of 40 accepted papers.
- NCSU Faculty Research and Professional Development Award, 2008.
- Second Invention Plateau Award, IBM T.J. Watson Research, 2007.
- CRA-W Distinguished Lecture, 2007.
- First Invention Plateau Award, IBM T.J. Watson Research, 2006.
- Invention Achievement Award, IBM T.J. Watson Research, 2004
- Saburo Muroga Fellowship to outstanding graduate students, 2003, UIUC
- ILLIAC Fellowship to best PhD candidates, 1999 - 2002, UIUC
- David J. Kuck Best Master Thesis Award, 2001, UIUC
- Motorola Scholarship to best undergraduate students, 1998, Peking University
- Outstanding undergraduate scholarships, 1995 - 1998, Peking University

Grants and Gifts (\$2.77 million in total)

- NSF CAREER Award “CAREER: Enabling Robust Virtualized Hosting Infrastructures via Coordinated Learning, Recovery, and Diagnosis”, \$450,000, 2012-2017, Sole PI.

- NSA Science of Security Lablet Grant “Deepening the Understanding of Least Privilege Through Automatic Partitioning of Hybrid Programs”, \$522K, 2012-2015, Co-PI, PI: William Enck.
- NSF CSR Grant “Online System Anomaly Prediction and Diagnosis for Large-Scale Hosting Infrastructures”, \$405,000, 2009-2012, Sole PI.
- NSF CSR Grant “Hybrid Opportunistic Computing for Green Clouds”, \$470,000, 2009-2012, Co-PI.
- NSF II:NEW:ARC: A Root Cluster for Systems Research in Scalable Computing, \$549,999, 2010-2013, Co-PI.
- Army Research Office Grant “Predictive Anomaly Management for Resilient Virtualized Computing Infrastructures”, \$300,000, 2010-2013, Sole PI.
- IBM Faculty Award 2011, \$15,000, Sole PI.
- Google Research Award 2011, \$61,000, Sole PI.
- IBM Faculty Award 2010, \$14,000, Sole PI.
- Google Research Award 2009, \$51,000, Sole PI.
- IBM Faculty Award, 2009, \$29,000, Sole PI.
- IBM Exploratory Stream Analytics Award, 2009-2010, \$20,000, Sole PI.
- IBM Faculty Award. 2008, \$30,000, Sole PI
- North Carolina CACC Grant, “Virtual Machine Management for Large-Scale Virtualized Hosting Infrastructure”, 2008-2009, \$40,000, Sole PI.
- NCSU FRPD award. Automatic 24x7 System Health Management for Large-Scale Hosting Infrastructure . 2006, \$6,000, Sole PI.

Conference and Workshop Publications (total number: 43)

1. Yongmin Tan, Hiep Nguyen, Zhiming Shen, Xiaohui Gu, Chitra Venkatramani, Deepak Rajan, “PREPARE: Predictive Performance Anomaly Prevention for Virtualized Cloud Systems, Proc. of International Conference on Distributed Computing Systems (**ICDCS**), Macau, China, June, 2012 (acceptance rate: 71/515=13.8%).
2. Zhiming Shen, Sethuraman Subbiah, Xiaohui Gu, John Wilkes, “CloudScale: Elastic Resource Scaling for Multi-Tenant Cloud Systems, Proc. of ACM Symposium on Cloud Computing (**SOCC**) in conjunction with **SOSP**, Cascais, Portugal, October, 2011 (acceptance rate: 30/178 = 16.8%).
3. Hiep Nguyen and Yongmin Tan and Xiaohui Gu, “Propagation-aware Anomaly Localization for Cloud Hosted Distributed Applications, Proc. of ACM Workshop on Managing Large-Scale Systems via the Analysis of System Logs and the Application of Machine Learning Techniques (**SLAML**) in conjunction with **SOSP**, Cascais, Portugal, October, 2011.
4. Kamal Kc, Xiaohui Gu, “ELT: Efficient Log-based Troubleshooting System for Cloud Computing Infrastructures, Proc. of IEEE International Symposium on Reliable Distributed Systems (**SRDS**), Madrid, Spain, October, 2011. (acceptance rate: 30/88 = 34%)
5. Juan Du, Nidhi Shah, Xiaohui Gu, “Adaptive Data-Driven Service Integrity Attestation for Multi-Tenant Cloud Systems”, Proc. of ACM/IEEE International Workshop on Quality of Service (**IWQoS**), San Jose, CA, June, 2011.(acceptance rate: 23/80 = 28%)

6. Yongmin Tan, Vinay Venkatesh, Xiaohui Gu, “OLIC: OnLine Information Compression for Scalable Hosting Infrastructure Monitoring”, Proc. of ACM/IEEE International Workshop on Quality of Service (**IWQoS**), San Jose, CA, June, 2011.(acceptance rate: 23/80 =28%)
7. Zhenhuan Gong, Xiaohui Gu, John Wilkes, “PRESS: PRedictive Elastic ReSource Scaling for Cloud Systems”, Proc. of IEEE International Conference on Network and Services Management (**CNSM**), Niagara Falls, Canada, October, 2010.(**Best Paper Award**, acceptance rate: 27/176 = 15%)
8. Juan Du, Xiaohui Gu, Ting Yu, “On Verifying Stateful Dataflow Processing Services in Large-Scale Cloud Systems”, Proc. of ACM Conference on Computer and Communications Security (**CCS**), poster session, Chicago, IL, October, 2010.
9. Zhenhuan Gong, Xiaohui Gu, “PAC: Pattern-driven Application Consolidation for Efficient Cloud Computing”, Proc. of IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (**MASCOTS**), Miami Beach, Florida, August, 2010. (**Best paper nominee**, acceptance rate: 16/116= 31%)
10. Yongmin Tan, Xiaohui Gu, “On Predictability of System Anomalies in Real World”, Proc. of IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (**MASCOTS**), Miami Beach, Florida, August, 2010. (acceptance rate: 16/116= 31%)
11. Yongmin Tan, Xiaohui Gu, Haixun Wang, “Adaptive Runtime Anomaly Prediction for Dynamic Hosting Infrastructures”, Proc. of ACM Symposium on Principles of Distributed Computing (**PODC**), Zurich, Switzerland, July, 2010. (acceptance rate: 21%)
12. Juan Du, Xiaohui Gu, Douglas S. Reeves, ”Highly Available Component Sharing in Large-Scale Multi-Tenant Cloud Systems”, Proc. of ACM International Symposium on High Performance Distributed Computing (**HPDC**), Chicago, IL, June, 2010. (acceptance rate: 25%)
13. Juan Du, Wei Wei, Xiaohui Gu, Ting Yu, “RunTest: Assuring Integrity of Dataflow Processing in Cloud Computing Infrastructures”, Proc. of ACM Symposium on Information, Computer and Communications Security (**ASIACCS**), Beijing, China, April, 2010. (acceptance rate: 15%)
14. Wei Wei, Juan Du, Ting Yu, Xiaohui Gu, “SecureMR: A Service Integrity Assurance Framework for MapReduce”, Proc. of Annual Computer Security Applications Conference (**ACSAC**), Honolulu, HI, December, 2009. (acceptance rate: 19%)
15. Juan Du, Wei Wei, Xiaohui Gu, Ting Yu, “Toward Secure Dataflow Processing in Open Distributed Systems”, Proc. of ACM Scalable Trusted Computing Workshop (STC) colocated with ACM Conference on Computer and Communications Security (**CCS**), Chicago, IL, November 2009 .
16. Zhenhuan Gong, Prakash Ramaswamy, Xiaohui Gu, Xiaosong Ma, “SigLM: Signature-Driven Load Management for Cloud Computing Infrastructures”. Proc. of IEEE International Conference on Quality of Service (**IWQoS**), Charleston, South Carolina, July, 2009.

17. Ying Zhao, Yongmin Tan, Zhenhuan Gong, Xiaohui Gu, Mike Wamboldt, "Self-Correlating Predictive Information Tracking for Large-Scale Production Systems". Proc. of IEEE International Conference on Autonomic Computing and Communications (**ICAC**), Barcelona, Spain, June, 2009. (acceptance rate: 15%)
18. Xiaohui Gu, Haixun Wang, "Online Anomaly Prediction for Robust Cluster Systems". Proc. of 25th International Conference on Data Engineering (**ICDE**), April, 2009. (full paper, acceptance rate: 17%)
19. Xiaohui Gu, Spiros Papadimitriou, Philip S. Yu, Shu-Ping Chang, "Toward Predictive Failure Management for Distributed Stream Processing Systems". Proc. of 24th International Conference on Distributed Computing Systems (**ICDCS**), June, 2008. (acceptance rate:16%)
20. Xiaohui Gu, Spiros Papadimitriou, Philip S. Yu, Shu-Ping Chang, "Online Failure Forecast for Fault-Tolerant Data Stream Processing". Proc. of 24th International Conference on Data Engineering (**ICDE**), April, 2008. (short paper)
21. Chen Chen, Xifeng Yan, Philip S. Yu, Jiawei Han, Dongqing Zhang, Xiaohui Gu, "Toward Graph Containment Search and Indexing", Proc. of Very Large Data Bases Conference (**VLDB**), 2007. (acceptance rate:16%)
22. Kun-Lung Wu, Philip S. Yu, Bugra Gedik, Kirsten W. Hildrum, Charu C. Aggarwal, Eric Bouillet, Wei Fan, David A. George, Xiaohui Gu, Gang Luo, Haixun Wang, "Challenges and Experience in Prototyping a Multi-Modal Stream Analytic and Monitoring Application on System S", Proc. of Very Large Data Bases Conference (**VLDB**), 2007.
23. Xiaohui Gu, Philip S. Yu, "Toward Self-Managed Media Stream Processing Service Overlays", In Proc. of IEEE International Conference on Multimedia and Expo (**ICME**), Beijing, China, July, 2007.
24. Jin Liang, Xiaohui Gu, Klara Nahrstedt, "Self-Configuring Information Management for Large-Scale Service Overlays", In Proc. of IEEE **INFOCOM**, Alaska, May, 2007. (acceptance rate:18%)
25. Xiaohui Gu, Zhen Wen, Philip S. Yu, "BridgeNet: An Adaptive Multi-Source Stream Dissemination Service Overlay", In Proc. of IEEE **INFOCOM Mini-Symposium**, Anchorage, AL, May, 2007. (acceptance rate: 25%)
26. Xiaohui Gu, Philip S. Yu, Haixun Wang, "Adaptive Load Diffusion for Multiway Windowed Stream Joins", In Proc. of IEEE International Conference on Data Engineering (**ICDE**), Istanbul, Turkey, April, 2007. (full paper, acceptance rate:18%)
27. Xiaohui Gu, Zhen Wen, Ching-Yung Lin, Philip S. Yu, "ViCo: An Adaptive Distributed Video Correlation System", Proc. of ACM Multimedia (**SIGMM**), Santa Barbara, CA, October, 2006. (full paper, acceptance rate: 17%)
28. Thomas Repantis, Xiaohui Gu, Vana Kalogeraki, "Synergy: Sharing-Aware Component Composition for Distributed Stream Processing Systems", Proc. of the 7th ACM/IFIP/USENIX International Middleware Conference (**Middleware**), Melbourne, Australia, November 2006 (acceptance rate: 17%)

29. Xiaohui Gu, Philip S. Yu, Klara Nahrstedt, "Optimal Component Composition for Scalable Stream Processing", Proc. of IEEE International Conference on Distributed Computing Systems (**ICDCS 2005**), Columbus, Ohio, June 6-10, 2005. (acceptance rate: 14%)
30. Xiaohui Gu, Philip S. Yu, "Adaptive Load Diffusion for Stream Joins", Proc. of ACM/IFIP/USENIX International Middleware Conference (**Middleware**), Grenoble, France, 2005. (acceptance rate: 21%)
31. Xiaohui Gu, Zhen Wen, Philip S. Yu, Zon-Yin Shae, "Supporting Multi-Party Voice-Over-IP Services with Peer-to-Peer Stream Processing", Proc. of ACM Multimedia (**SIGMM**), Singapore, November, 2005. (short paper)
32. Xiaohui Gu, Klara Nahrstedt, Bin Yu, "An Integrated Peer-to-Peer Service Composition Framework ", Proc. of IEEE International Symposium on High-Performance Distributed Computing (**HPDC**), Honolulu, Hawaii, June 4-6, 2004. (acceptance rate: 15%)
33. Xiaohui Gu, Klara Nahrstedt, Rong N. Chang, Zon-Yin Shae, "An Overlay Based QoS-Aware Voice-Over-IP Conferencing System", in Proc. of IEEE International Conference on Multimedia and Expo 2004 (**ICME**). (acceptance rate: 30%)
34. Xiaohui Gu, Klara Nahrstedt, Rong N. Chang, Christopher Ward, "QoS-Assured Service Composition in Managed Service Overlay Networks", in Proc. of IEEE International Conference on Distributed Computing Systems (**ICDCS**), Providence, Rhode Island, May 19-22, 2003. (acceptance rate: 17%)
35. Xiaohui Gu, Klara Nahrstedt, Alan Messer, Ira Greenberg, Dejan Milojicic, "Adaptive Offloading Inference for Delivering Applications in Pervasive Computing Environment", Proc. of IEEE International Conference on Pervasive Computing and Communications (**PerCom**), Dallas-Fort Worth, Texas, March 23-26, 2003. (acceptance rate: 30%)
36. Duangdao Wichadakul, Xiaohui Gu, Klara Nahrstedt, "A Programming Framework for Ubiquitous Multimedia Applications", Proc. of ACM Multimedia 2002 (**SIGMM**), Juan Les Pins, France, December 1-6, 2002. (full paper, acceptance rate: 15%)
37. Xiaohui Gu, Klara Nahrstedt, "A Scalable QoS-Aware Service Aggregation Model for Peer-to-Peer Computing Grids", Proc. of 11th IEEE International Symposium on High Performance Distributed Computing (**HPDC**), Edinburgh, Scotland, July 24-26, 2002. (acceptance rate: 30%)
38. Xiaohui Gu, Klara Nahrstedt, "Dynamic QoS-Aware Multimedia Service Configuration in Ubiquitous Computing Environments", Proc. of IEEE International Conference on Distributed Computing Systems (**ICDCS**), Vienna, Austria, July 2-5, 2002. (acceptance rate: 18%)
39. A. Messer, I. Greenberg, P. Bernadat, D. Milojicic, D. Chen, T.J. Guili, X. Gu, "Towards a Distributed Platform for Resource-Constrained Devices", Proc. of IEEE International Conference on Distributed Computing Systems (**ICDCS**), Vienna, Austria, July 2-5, 2002. (acceptance rate: 18%)

40. Duangdao Wichadakul, Klara Nahrstedt, Xiaohui Gu, Dongyan Xu, “2KQ+: An Integrated Approach of QoS Compilation and Reconfigurable, Component-based Runtime Middleware for the Unified QoS Management Framework”, Proc. of IFIP/ACM International Conference on Distributed Systems Platforms (**Middleware**), Heidelberg, Germany, Nov. 2001 (acceptance rate: 17%)
41. Xiaohui Gu, Klara Nahrstedt, “An Event-Driven, User-Centric, QoS-aware Middleware Framework for Ubiquitous Multimedia Applications ”, Proc. of ACM Multimedia (Multimedia Middleware Workshop), Ottawa, Canada, September 30 - October 5, 2001.
42. Wanghong Yuan, Klara Nahrstedt, Xiaohui Gu, “Coordinating Energy-Aware Adaptation of Multimedia Applications and Hardware Resource,” Proc. of 9th ACM Multimedia (Multimedia Middleware Workshop), Ottawa, Canada, September 30 - October 5, 2001.
43. Xiaohui Gu, Duangdao Wichadakul, Klara Nahrstedt , “Visual QoS Programming Environment for Ubiquitous Multimedia Services”, Proc. of IEEE International Conference on Multimedia and Expo 2001(**ICME**), Tokyo, Japan, August 22 - August 25, 2001. (acceptance rate: 30%)

Journal Publications (total number: 7)

1. Thomas Repantis, Xiaohui Gu, Vana Kalogeraki, “QoS-Aware Shared Component Composition for Distributed Stream Processing Systems”, IEEE Transactions on Parallel and Distributed Systems (**TPDS**), 20(7): pp. 968-982, July 2009.
2. Xiaohui Gu, Zhen Wen, Philip S. Yu, Zon-Yin Shae, “peerTalk: A Peer-to-Peer Voice-Over-IP System”, IEEE Transactions on Parallel and Distributed Systems (**TPDS**), 19(3): 2008.
3. Xiaohui Gu, Klara Nahrstedt, “On Composing Stream Applications in Cooperative Peer-to-Peer Environments”, IEEE Transactions on Parallel and Distributed Systems (**TPDS**), 17(8): pp. 824-837, 2006.
4. Xiaohui Gu, Klara Nahrstedt, “Distributed Multimedia Service Composition with Statistical QoS Assurances”, IEEE Transactions on Multimedia (**TMM**), 8(1): pp. 141-151, 2006.
5. Xiaohui Gu, Alan Messer, Ira Greenberg, Dejan Milojicic, Klara Nahrstedt, “Adaptive Offloading for Pervasive Computing”, IEEE **Pervasive Computing Magazine**, Vol. 3, No. 3, July, 2004.
6. Xiaohui Gu, Klara Nahrstedt, “QoS-Aware Service Composition for Large-Scale Peer-to-Peer Systems”, Book Chapter in Grid Resource Management: State of the Art and Future Trends, Chapter 24, pp. 395-410, Kluwer Academic Press, October, 2003.
7. Xiaohui Gu, Klara Nahrstedt , Wanghong Yuan, Duangdao Wichadakul, Dongyan Xu, “An XML-based Quality of Service Enabling Language for the Web”, Journal of Visual Language and Computing(**JVLC**), special issue on multimedia languages for the Web, volume 13, No.1, pp. 61-95, Academic Press, February, 2002.

Refereed Posters and Work-in-Progress Presentations (total number: 4)

1. R. Benjamin Clay, Zhiming Shen, Xiaosong Ma, Xiaohui Gu, "Augmenting MapReduce with Active Volunteer Resources", Poster Session in ACM Symposium on Operating Systems Principles (**SOSP**), Cascais, Portugal, October, 2011.
2. Zhenhuan Gong, Xiaohui Gu, "Predictive Elastic Load Management for Cloud Computing Infrastructures", Poster Session in ACM Symposium on Operating Systems Principles (**SOSP**), Big Sky, MT, October, 2009.
3. Xiaohui Gu, Klara Nahrstedt, "Scalable Constrained Routing in Overlay Networks", Work In Progress Session in USENIX Operating Systems Design and Implementation (**OSDI**), Boston, MA, December 2002.
4. Xiaohui Gu, Klara Nahrstedt, "QoS Talk: QoS Programming Environment for Ubiquitous Multimedia Applications", Poster Session at ACM SIGPLAN Conference on Programming Language Design and Implementation (**PLDI**), Snowbird, UT, June 2001.

Patents (total number: 8)

1. Xiaohui Gu, Rong N. Chang, Zon-Yin Shae, System and Method for Providing a Communication Session, US Patent No. 7302053, Nov. 27th, 2007.
2. Xiaohui Gu, Spyridon Papadimitriou, Shu-Ping Chang, Philip Yu, System and Methods for Predictive Failure Management, Filed with the US patent office, 2007.
3. Haixun Wang, Xiaohui Gu, Lipyeow Lim, Min Wang, System and Method for Indexable XML Data Serialization, Filed with the US patent office, 2007.
4. Xiaohui Gu, Philip S. Yu, Shu-Ping Chang, Model-based Self-Optimizing Distributed Information Management, Filed with the US patent office, 2006.
5. Xiaohui Gu, Philip S. Yu, Haixun Wang, System and Method for Scalable Processing of Multi-way Data Stream Correlations, Filed with the US patent office, 2006.
6. Xiaohui Gu, Zhen Wen, Philip S. Yu, Zon-Yin Shae, System and Method for Peer-to-Peer Multi-Party Voice-Over-IP Services, Filed with the US patent office, 2006.
7. Xiaohui Gu, Philip S. Yu, Method and Apparatus for Providing Load Diffusion in Data Stream Correlations, Filed with the US patent office, 2005.
8. Xiaohui Gu, Philip S. Yu, System and Methods for Optimal Component Composition in a Stream Processing System, Filed with the US patent office, 2005.

Selected Professional Services

- **Recent Invited Talks:** CRA-W Distinguished Lecture at Kentucky Women in Computer Science Conference (February, 2012), Columbia University (July 2011), IBM Research (July, 2011), HP Labs (June, 2011), NEC Labs (January 2009), IBM Research (September, 2008)
- **Program Co-Chair:** IWQoS 2013
- **TPC Member:** ICDCS 2012-2011, 2009, Middleware 2012 (main research track), Middleware 2011 (industry track), IWQoS 2012, 2011, 2009, MDM 2011, 2009, ICME 2010-2009, BroadNets 2010, SDM 2009, IPDPS 2009, PerCom 2006-2009, SIGMM 2008, ICC 2008, Middleware (Doctoral Symposium) 2007, ICPDS 2007, MMM 2006, RTSS 2006, etc.
- **Journal Review:** IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Mobile Computing, IEEE Transactions on Networking, ACM Transactions on

Internet Technology , ACM Transactions on Multimedia,IEEE Transactions on Computers, IBM Systems Journal, Springer Journal of Internet Services and Applications, Special Issue on Cloud Computing, IEEE Internet Computing, etc.

- **Outreach activity:** Faculty advisor for Women in Computer Science (WiCS) at NCSU, Host for CRA-W Distinguished Lecture 2011.
- **Proposal Review:** NSF Review Panelist, 2010, 2012, External Reviewer for ERC European Research Council, 2010, External Reviewer for Research Grants Council (RGC) of Hong Kong, 2007.
- Senior Member of IEEE and member of ACM.

Current Students

- Yongmin Tan (PhD, pass prelim exam in Nov. 2011, expected to graduate in Summer 2012)
- Kamal Kc (PhD, pass qualify exam in fall 2010)
- Zhiming Shen (PhD, Fall 2010-present)
- Hiep Nyugen (PhD, Summer 2010-present)
- Daniel Dean (PhD, Spring 2011-present)
- Anson Ho (PhD, Fall 2011-present)

Graduated Students

- Juan Du (PhD, 2011, Amazon)
- Sethuraman Subbiah (MS, 2011, NetApp Advanced Technology Group (ATG))
- Vinay Venkatesh (MS 2010, IBM RTP)
- Prakash Ramaswamy (MS 2009, EDL consulting, Chicago)

Miscellaneous

Female, US Permanent Resident