

MING ZHAO

Advanced Computing and Information System Lab (ACIS)
Electrical and Computer Engineering, University of Florida
P.O. Box 116200, 339 Larsen Hall, Gainesville, FL 32611

Phone: (352) 870-2036, Fax: (352) 392-5040
Email: ming@acis.ufl.edu
Web: www.acis.ufl.edu/~ming

SUMMARY OF QUALIFICATIONS

- Extensive research experience in computer systems and a solid background in computer science and engineering, control, and machine learning
- Actively published in leading journals and conferences in distributed and grid computing, autonomic computing, such as HPDC, ICDCS, ICAC, SC; and contributed valuable software to the community
- Substantial interdisciplinary research in several large projects in collaboration with experts from biomedical science, coastal engineering, signal processing, and neuroscience
- Excellent speaking and communication skills developed through conference presentations and guest lectures; teaching experience in assisting both undergraduate and graduate courses

RESEARCH INTERESTS

Distributed systems, virtualization, grid and high-performance computing, autonomic computing, cyberinfrastructure, file and storage systems, and operating systems

EDUCATION

- Ph.D. in Electrical and Computer Engineering** May 2008
University of Florida, Gainesville, FL
Dissertation: File System Virtualization and Autonomic Services for Grid Data Management
Committee: Renato J. Figueiredo (Chair), José A. B. Fortes, Alan D. George, Sanjay Ranka
- M.S. in Pattern Recognition and Intelligent Systems** July 2001
Tsinghua University, Beijing, China
Dissertation: A Human-computer Interaction System for Content-based Image Retrieval
Committee: Gang Rong (Chair), Zhaoqi Bian, Changshui Zhang, Jie Zhou
- B.E. in Automation**, Excellent Graduate Award winner July 1999
Tsinghua University, Beijing, China
Dissertation: An Approach to Automated Palmprint Recognition Based on K-L Transform

RESEARCH EXPERIENCE

Research Assistant January 2003 - Present
Advanced Computing and Information Systems Lab, University of Florida, Gainesville, FL
Advisor: Renato J. Figueiredo

- **User-level File System Virtualization and Enhancements**
Designed and built GVFS, a wide-area virtual file system that supports transparent cross-domain data access for unmodified applications and operating systems, and provides application-tailored enhancements for performance, consistency, security, and reliability in dynamic, heterogeneous, and large-scale environments [C3][C1][I1]. (*URL*: www.acis.ufl.edu/~ming/gvfs)

- **Service-based Autonomic Grid Data Management**
 Proposed a service-based architecture for autonomic management of grid-wide data provisioning [C4][C5]. Developed Web Service Resource Framework based middleware to automatically schedule and configure dynamic data sessions, in accordance with application-desired service quality.
 (URL: www.acis.ufl.edu/~ming/dms)
- **Cyberinfrastructure for Dynamic Data-driven Brain-machine Interfaces (BMI)**
 Developed a middleware-based cyberinfrastructure that consolidates distributed resources to support the time-critical and resource-demanding parallel BMI computing. Also provided scientists with a user-friendly and powerful Web-based portal to manage online and offline experiments [C8][I2].
 (URL: www.acis.ufl.edu/dddasbmi)
- **QoS-driven Resource Management of Virtualized Data Centers**
 Designed and built an autonomic system for managing virtualized resources based on application workloads and service quality targets, which supports both fine-grained resource allocation using machine-learning-based prediction and a profit-driven model [J3][C7], and advance resource reservation leveraging VM migrations to transparently vacate workloads [C2][T1].
 (URL: www.acis.ufl.edu/~ming/vrm)
- **On-demand Virtual Machine (VM) Environments for Grid Computing**
 Designed and implemented a data management solution to provide high-performance and secure access to distributed VM states, which allows fast instantiation and efficient execution of VMs to provide flexible, customizable, and secure execution environments for grid computing [J1][C6].
 (URL: www.acis.ufl.edu/~ming/vmgc)
- **In-VIGO Virtualization Middleware for Computational Grids**
 Developed and maintained the data management component for In-VIGO, a grid middleware system that employs virtualization extensively, from virtual data to virtual applications [J2][C9]. In-VIGO has been in production since 2003 and has supported thousands of scientific job executions, including biomedical study and costal engineering modeling [C10]. (URL: invigo.acis.ufl.edu)

Research Intern

May 2006 - August 2006

IBM Research Division, Austin Research Lab, Austin, TX
 Mentor: Eric Van Hensbergen

- **Dynamic Policy Disk Caching for Storage Networking**
 Proposed client-side disk caching to solve the scalability problem for storage networking [T2]. Designed and implemented the dm-cache kernel module, a generic block-level disk cache that supports transparent caching for different types of storage networks and enables policy-guided dynamic cache customizations. (URL: www.acis.ufl.edu/~ming/dmcache)

Research Assistant

August 1998 - July 2001

National Key Lab of Intelligent Technology and Systems, Tsinghua University, Beijing, China
 Advisor: Gang Rong

- **A Human-computer Interaction System for Content-based Image Retrieval (CBIR)**
 Proposed a novel algorithm that combines various image features (color and texture) and user feedback for CBIR, and implemented a Web-based system to allow interactions with users and leverage their feedback on the retrieved images to automatically refine the retrieval iteratively [J4].
- **An Approach to Automated Palmprint Recognition**
 Developed a dual-eigenspace algorithm for automated palmprint recognition by using the Karhunen-Loeve transform to extract features (eigenpalms) from training set palmprints and classifying a palmprint with Euclidean distance in the subspace spanned by the eigenpalms [J6].

TEACHING EXPERIENCE

Teaching Assistant, Guest Lecturer Fall 2006, Fall 2007
EEL6892 - Virtual Computers, University of Florida
 Prof. Renato J. Figueiredo

Instructed lectures on virtual storage. Topics included concepts and principles of storage virtualization and services, as well as case studies and demonstrations of advanced virtual storage systems and their integration with virtual machines. (*Lectures URL: www.acis.ufl.edu/~ming/lectures*)

Teaching Assistant Spring 2001
UNIX Operating System, Tsinghua University
 Prof. Huirong Chen

Instructed and guided students in lab assignments. Graded homework, projects, and exams.

PUBLICATIONS**Journals**

- [J1] **M. Zhao**, J. Zhang, and R. J. Figueiredo, "Distributed File System Virtualization Techniques Supporting On-demand Virtual Machine Environments for Grid Computing," *Journal of Cluster Computing*, vol. 9, no. 1, pp. 45-56, January 2006.
- [J2] S. Adabala, V. Chadha, P. Chawla, R. Figueiredo, J. Fortes, I. Krsul, A. Matsunaga, M. Tsugawa, J. Zhang, **M. Zhao**, L. Zhu, and X. Zhu, "From Virtualized Resources to Virtual Computing Grids: the In-VIGO System," *Journal of Future Generation Computing Systems*, vol. 21, no. 6, pp. 896-909, June 2005.
- [J3] J. Xu, **M. Zhao**, J. A. B. Fortes, R. Carpenter, and M. Yousif, "Autonomic Resource Management in Virtualized Data Centers using Fuzzy Logic-based Approaches," (to appear) *Journal of Cluster Computing*.
- [J4] **M. Zhao** and G. Rong, "A Human-computer Interaction System for Content-based Image Retrieval on Internet," *Journal of Tsinghua University (Science and Technology, Special Issue)*, August 2001.
- [J5] Z. Wang, **M. Zhao**, and Q. Yu, "Modeling of Branching Structures of Plants," *Journal of Theoretical Biology*, vol. 209, no. 4, pp. 383-394, April 2001.
- [J6] **M. Zhao**, W. Shu, G. Rong, and Z. Bian, "An Approach to Automated Palmprint Recognition Based on K-L Transform," *Journal of Tsinghua University (English abstract available on Ei Compendex)*, vol. 40, no. 9, pp. 100-103, September 2000.

Conference Proceedings

- [C1] **M. Zhao** and R. J. Figueiredo, "A User-level Secure Grid File System," *Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis (SC07)*, November 2007. (**Acceptance Rate: 20%**)
- [C2] **M. Zhao** and R. J. Figueiredo, "Experimental Study of Virtual Machine Migration in Support of Reservation of Cluster Resources," *Proceedings of 2nd International Workshop on Virtualization Technology in Distributed Computing (held in conjunction with SC07) (VTDC 2007)*, November 2007.
- [C3] **M. Zhao** and R. J. Figueiredo, "Application-tailored Cache Consistency for Wide-area File Systems," *Proceedings of 26th International Conference on Distributed Computing Systems (ICDCS 2006)*, July 2006. (**Acceptance Rate: 13%**)

- [C4] **M. Zhao**, J. Xu, and R. J. Figueiredo, "Towards Autonomic Grid Data Management with Virtualized Distributed File Systems," *Proceedings of 3rd IEEE International Conference on Autonomic Computing (ICAC 2006)*, pp. 209-218, June 2006. (**Acceptance Rate: 20%**)
- [C5] **M. Zhao**, V. Chadha, and R. J. Figueiredo, "Supporting Application-tailored Grid File System Sessions with WSRF-based Services," *Proceedings of 14th IEEE International Symposium on High Performance Distributed Computing (HPDC 2005)*, pp. 24-33, July 2005. (**Top 5 Best Papers, Acceptance Rate: 17%**)
- [C6] **M. Zhao**, J. Zhang, and R. J. Figueiredo, "Distributed File System Support for Virtual Machines in Grid Computing," *Proceedings of 13th IEEE International Symposium on High Performance Distributed Computing (HPDC 2004)*, pp. 202-211, June 2004. (**Acceptance Rate: 15%**)
- [C7] J. Xu, **M. Zhao**, J. A. B. Fortes, R. Carpenter, and M. Yousif, "On the Use of Fuzzy Modeling in Virtualized Data Center Management," *Proceedings of 4th International Conference on Autonomic Computing (ICAC 2007)*, June 2007. (**Best Student Paper, Acceptance Rate: 14%**)
- [C8] J. DiGiovanna, L. Marchal, P. Rattanamrong, **M. Zhao**, S. Darmanjian, B. Mahmoudi, J. C. Sanchez, J. C. Principe, L. Hermer-Vazquez, R. Figueiredo, and J. Fortes, "Towards Real-time Distributed Signal Modeling for Brain Machine Interfaces," *Proceedings of International Conference on Computational Science*, pp. 964-971, May 2007.
- [C9] A. Matsunaga, M. Tsugawa, **M. Zhao**, L. Zhu, V. Sanjeevan, S. Adabala, R. Figueiredo, H. Lam, and J. A. B. Fortes, "On the Use of Virtualization and Service Technologies to Enable Grid-computing," *Proceedings of 11th International Euro-Par Conference*, pp. 1-12, August 2005.
- [C10] J. Paladugula, **M. Zhao**, and R. J. Figueiredo, "Support for Data-intensive, Variable-granularity Grid Applications via Distributed File System Virtualization—A Case Study of Light Scattering Spectroscopy," *Proceedings of Challenges of Large Applications in Distributed Environments (held in conjunction with HPDC 2004) (CLADE/2004)*, pp. 12-21, June 2004.

Technical Reports

- [T1] **M. Zhao** and R. J. Figueiredo, "Virtual Machine Based Resource Reservation for Real-time Tasks," *Technical Report (TR-ACIS-07-01)*, University of Florida, October 2007.
- [T2] E. V. Hensbergen and **M. Zhao**, "Dynamic Policy Disk Caching for Storage Networking," *IBM Research Report (RC24123)*, November 2006.
- [T3] **M. Zhao** and R. J. Figueiredo, "Proxy Managed Client-side Disk Caching for the Virtual File System," *Technical Report (TR-ACIS-04-01)*, University of Florida, November 2004.

Working Papers

- [I1] **M. Zhao** and R. J. Figueiredo, "Application-tailored User-level Wide-area File Systems," submitted to *IEEE Transactions on Parallel and Distributed Systems*.
- [I2] **M. Zhao**, P. Rattanamrong, R. J. Figueiredo, and J. A. B. Fortes, "A Cyber-workstation for Dynamic Data-driven Brain-machine Interfaces," in preparation for conference submission.
- [I3] J. Xu, **M. Zhao**, and J. A. B. Fortes, "Cooperative Autonomic Management in Dynamic Distributed Systems," in preparation for journal submission.

PRESENTATIONS

- [P1] "A User-level Secure Grid File System," *Supercomputing*, November 2007.
- [P2] "Experimental Study of Virtual Machine Migration in Support of Reservation of Cluster Resources," *VTDC*, November 2007.
- [P3] "Study of Erasure Coding and its Application on Building Reliable Distributed File System," *ECE, Univ. of Florida*, April 2007.
- [P4] "Dynamic Policy Disk Cache for Storage Networking," *IBM Research Seminar*, August 2006.
- [P5] "Application-tailored Cache Consistency for Wide-area File Systems," *ICDCS*, July 2006.
- [P6] "Towards Autonomic Grid Data Management with Virtualized Distributed File Systems," *ICAC*, June 2006.
- [P7] "Performance Analysis of Xen: Intel VT-x vs. Native x86," *ECE, Univ. of Florida*, April 2006.
- [P8] "Supporting Application-tailored Grid File System Sessions with WSRF-based Services," *HPDC*, July 2005.
- [P9] "A Case Study of Using MPI on Clusters in Grid Environments," *ECE, Univ. of Florida*, April 2005.
- [P10] "Distributed File System Support for Virtual Machines in Grid Computing," *HPDC*, June 2004.
- [P11] "Support for Data-intensive, Variable-granularity Grid Applications via Distributed File System Virtualization—A Case Study of Light Scattering Spectroscopy," *CLADE*, June 2004.
- [P12] "Layer 3 Virtual Networking for Virtual Machines," *ECE, Univ. of Florida*, December 2003.

SOFTWARE CONTRIBUTIONS

GVFS (Grid Virtual File System)

A user-level virtual file system for transparent grid-wide data access with application-tailored enhancements on caching, consistency, security, and fault-tolerance. Implemented as user-level Network File System proxies, and developed in C. ([URL: www.acis.ufl.edu/~ming/gvfs](http://www.acis.ufl.edu/~ming/gvfs))

DMS (Data Management Services)

Web services for managing data provisioning based on GVFS, Secure Copy, and GridFTP. Implemented in Perl and Java. Built upon Web Service Resource Framework with WSRF::Lite and Globus Toolkit 4. ([URL: www.acis.ufl.edu/~ming/dms](http://www.acis.ufl.edu/~ming/dms))

DDDAS-BMI Cyber-workstation

A distributed software system that provides in vivo data acquisition, real-time robot control, reliable messaging, parallel BMI model computing, and an AJAX-portlet based Web management interface. Implemented in C++ (Windows and Linux), Java, and JavaScripts. ([URL: bmi.acis.ufl.edu](http://bmi.acis.ufl.edu))

DM-Cache

A Linux kernel module for generic and highly-customizable block-level disk caching. Implemented in C, and built upon the Linux device-mapper framework. ([URL: www.acis.ufl.edu/~ming/dmcache](http://www.acis.ufl.edu/~ming/dmcache))

SRPC (Secure Remote Procedure Call) Library

A generic library for secure remote procedure calls. Implemented in C. Built upon OpenSSL and TI-RPC to provide strong security and transport independence. ([URL: www.acis.ufl.edu/~ming/srpc](http://www.acis.ufl.edu/~ming/srpc))

PROFESSIONAL SERVICE

Web Chair of 5th IEEE International Conference on Autonomic Computing (2008)

Web Chair of 4th IEEE International Conference on Autonomic Computing (2007)

Official Journal Referee

- IEEE Transactions on Parallel and Distributed Systems (2006)

External Conference Reviewer

- IEEE International Symposium on High Performance Distributed Computing (2004 - 2008)
- IFIP International Conference on Network and Parallel Computing (2007)
- International Conference on Distributed Computing Systems (2006)
- IEEE International Conference on Autonomic Computing (2006)
- International Meeting on High Performance Computing for Computational Science (2006)
- International Workshop on Virtualization Technology in Distributed Computing (2006)
- International Conference on High Performance Computing (2004)

Student Volunteer of International Conference for High Performance Computing, Networking, Storage and Analysis (2007), and IEEE International Conference on Autonomic Computing (2007)

Member of IEEE (Computer Society) and ACM

Vice President of FACSS (Friendship Association of Chinese Students and Scholars) of the University of Florida (2003 - 2004)

COMPUTER SKILLS

- Parallel programming, Grid programming, Web services, Network simulation and programming, Kernel development, Security programming, Database programming, Digital circuit design etc.
- **Languages:** C/C++, Java, Perl, SQL, Assembly, HTML/XML, JavaScript, ASP/JSP/PHP etc.
- **Platforms:** Linux, UNIX, Windows, Virtual Machines (Xen, VMware, UML, z/VM) etc.
- **Tools:** MPI, Globus Toolkit, MATLAB, NS2, Cadence, Simple Scalar etc.

RELEVANT COURSES

- Virtual Computers, Distributed Computing, Operating Systems and Concurrency Programming
- Billion Transistor Computer Architecture, Fault Tolerant Computer Architecture, Parallel Computer Architecture, Computer Systems Architecture
- Mobile and Wireless Networking, Computer Communications, Computer Networks
- Nano Computing, Advanced VLSI Design, VLSI Circuits and Technology
- Programming Language Principles, Analysis of Algorithms, Database Systems
- Artificial Intelligence, Pattern Recognition, Image Analysis and Computer Vision, Artificial Neural Networks, Cognitive Science
- Automatic Control Theory, Process Control, System Identification, Computer Control Systems
- Numerical Analysis and Algorithms, Probability and Statistics, Operation Research