Yang Wang

Department of Computer Science
The University of Texas at Austin
1 University Station C0500
Austin, Texas 78712-0233 USA

Office: ACES 6SEo4D
phone: (512)695-9122
email: yangwang@cs.utexas.edu
http: //www.cs.utexas.edu/~yangwang

Research Interests

Fault tolerance in large-scale distributed systems

Education

Ph.D. in Computer Science — The University of Texas at Austin, since 2008
Advisor: Prof. Mike Dahlin

M.E. in Computer Science & Technology — Tsinghua University, 2005-2008
Advisor: Prof. Jiwu Shu *Topic*: Adaptive Caching Mechanism

B.E. in Computer Science & Technology — Tsinghua University, 2001-2005

Honors and Awards

2006 Second class scholarship of McData.
2004 First class scholarship of Tsinghua University.
2001-2003 Second class scholarship of Tsinghua University.

Research and Working Experience

2008-Present UpRight Cluster Services — The University of Texas at Austin
Advisor: Prof. Mike Dahlin and Lorenzo Alvisi
This work aims to make Byzantine Fault Tolerance (BFT) a practical alternative to build high-availability distributed systems. It significantly reduces the efforts required to modify an application, while still maintains an acceptable performance. My work focuses on the application side, including: 1) define a clear and easy to understand interface between the BFT module and the application; 2) modify Hadoop Zookeeper Distributed Lock Service to incorporate it into the UpRight framework. The major challenge is to eliminate nondeterminisms in the application; 3) design and implement a general and efficient checkpoint and recovery mechanism for the applications.

2007-2008 Tracing File System Operations in Real Environments — Tsinghua University
Advisor: Prof. Jiwu Shu
We designed a tool to trace file system operations in real environments for file system evaluation and I/O performance optimization. The major challenge comes from the trace provider, which does not allow us to stop the applications to install this tool. This tool can be loaded dynamically and then injected between the virtual file system (VFS) layer and the applications. It has a low-overhead compressing mechanism to reduce the size of traces, thus to reduce I/Os. The runtime overhead is about 5%, making it acceptable in real environments. It can support different kinds of file systems on Linux. I led the design, implementation, and evaluation of this project.

2005-2007 Selecting the Optimal Policy Adaptively for a cache system — Tsinghua University
Advisor: Prof. Jiwu Shu
We designed a cache system which can switch to different caching policies based on the observation of current workload. First, we design a general interface so that we can encapsulate caching policies into modules and switch them between each other. Second, we log cache access traces online and analyze them offline to find out the optimal policy. The prototype evaluation demonstrated that it overcomes the well-known adaptive policy by up to 11.9% in average response time. I led the design, implementation and evaluation of this project.

2006-2007 Large scale random accessible storage system Netease R&D (Internship)
I worked on a project of building a terabyte scale random accessible storage system over hundreds of machines. It stores half-structured data and supports both random search and sequential scan operations. I am one of the core developers and maintainers. The system is now used to support one of the largest search engines in China.

2005-2006 Out-of-band storage virtualization in Storage Area Network Tsinghua University
Advisor: Prof. Jiwu Shu
We built a system that can aggregate multiple disks into volumes and share these volumes across many clients. The system uses out-of-band architecture, which separates data flow from control flow, to achieve high performance and scalability. It supports asymmetric platforms, including Windows, Linux and Solaris. The clients cache metadata to improve performance and reliability. I worked on the Solaris virtualization module during this project.

Teaching Experience

2009 Spring Teaching Assistant, UT Austin, course CS345 Programming Languages
2008 Fall Teaching Assistant, UT Austin, course CS352 Computer Systems Architecture

Publications

[1] Allen Clement, Manos Kapritsos, Sangmin Lee, Yang Wang, Lorenzo Alvisi, Mike Dahlin, T. Riche, UpRight Cluster Services, *Proceedings of the 22 nd ACM Symposium on Operating Systems Principles (SOSP)*, Oct 2009 .

[2] Yang Wang, Jiwu Shu, Wei Xue, Mao Xue, VFS Interceptor: Dynamically Tracing File System Operations in real environments. *First International Workshop on Storage and I/O Virtualization, Performance, Energy, Evaluation and Dependability (SPEED2008)*. Held in conjunction with *HPCA'08*.

[3] Jiwu Shu, Yang Wang, Wei Xue, Yifeng Luo. An Efficient SAN-Level Caching Method Based on Chunk-Aging. *International Conference on Networking, Architecture and Storage 2007 (NAS 2007)*. Best Student Presentation Award.

[4] Yang Wang, Wei Xue, Ji-Wu Shu, Guang-Yan Zhang. Design and implementation of an out-of-band virtualization system on Solaris 10. *International Conference on Computational Science 2006 (ICCS 2006) IST Workshop*.

Extracurricular Activities

2004 Summer Trend Micro Programming Contest, Top 10 of 305 teams in mainland China. Best User Interface Award.