

Yi Qiao

yqiao@cs.northwestern.edu

http://www.cs.northwestern.edu/~yqiao

Department of Computer Science
Northwestern University
1890 Maple Avenue, Suite 300
Evanston, IL 60201

(847)491-7060 (Office)
(847)722-5862 (Cell)
(847)491-5258 (Fax)

Research Interests

I am broadly interested in distributed systems and networking, with an emphasis on the design, measurement, modeling, and performance analysis of more scalable, efficient, and resilient systems. Recent projects mainly focus on design and implementation of peer-to-peer/Grid systems, queuing and scheduling analysis, wide area traffic measurement and modeling. My research takes the approach that covers both experimental buildup, theoretical modeling, and performance analysis. My solid background in statistics, queuing theory, plus my system building and evaluation capability gained through several projects, could give me the power of delving deeper into many existing and potential intriguing territories in large-scale system research.

Education

Northwestern University, Evanston, IL, USA

Ph.D. Candidate in *Computer Science*

Expected June, 2007

Advisor: Fabian E. Bustamante

M.S. in *Computer Science*

December, 2003

Tsinghua University, Beijing, China

M.S. in *Electrical Engineering*

June, 2001

Tsinghua University, Beijing, China

B.S. in *Electrical Engineering*

June, 1998

Employment

Northwestern University, Department of Computer Science

Sept 2002 to present

Research Assistant for several projects.

Teaching Assistant twice for *CS343 Operating Systems* and *CS339 Introduction to Database Systems*.

Co-advising undergraduate students in their senior projects.

Tsinghua University, Department of Electrical Engineering

Sept 1998 to June 2001

Research Assistant for several projects.

Teaching Assistant once for *Dynamic Power System Analysis*.

Publications in Computer Science

Refereed Publications

- Dong Lu, Yi Qiao, Peter Dinda, Fabian Bustamante, “Characterizing and Predicting TCP Throughput on the Wide Area Network”, *Proc. of the 25th International Conference on Distributed Computing Systems (ICDCS 2005)*, June 2005, Columbus, Ohio. To appear.
- Dong Lu, Yi Qiao, Peter Dinda, Fabian Bustamante, “Modeling and Taming Parallel TCP on the Wide Area Network”, *Proc. of the 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2005)*, April 2005, Denver, Colorado. To appear.
- Yi Qiao, Dong Lu, Fabian E. Bustamante and Peter Dinda, “Looking at the Server-Side of Peer-to-Peer Systems”, *Proc. of the 7th Workshop on Languages, Compilers and Run-time Support for Scalable Systems (LCR 2004)*, October 2004, Houston, TX.
- Y. Qiao, J. Skicewicz, P. Dinda, “An Empirical Study of the Multiscale Predictability of Network Traffic”, *Proc. of the 13th IEEE International Symposium on High Performance Distributed Computing (HPDC 2004)*, June 2004, Honolulu, Hawaii.
- Fabian E. Bustamante and Yi Qiao, “Friendships that last: Peer lifespan and its role in P2P protocols”, *Proceedings of the Eighth International Workshop on Web Content Caching and Distribution (WCW 2003)*, Sept. 2003, Hawthorne, NY.
- Yi Qiao and Fabian E. Bustamante. “Elders Know Best: Lifespan-Based Ideas in P2P Systems”, *Work-In-Progress, 19th Symposium on Operating Systems Principles (SOSP 2003)*, October 2003, Bolton Landing, NY.
- Dong Lu, Peter Dinda, Yi Qiao, Huanyuan Sheng, Fabian Bustamante, “Applications of SRPT Scheduling with Inaccurate Scheduling Information”, *Proceedings of 12th Annual Meeting of the IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2004)*, October 2004, Volendam, The Netherlands.
- Stefan Birrer, Dong Lu, Fabian Bustamante, Yi Qiao, Peter Dinda, “FatNemo: Building a Resilient Multi-Source Multicast Fat-Tree”, *Proceedings of the Ninth International Workshop on Web Content Caching and Distribution (WCW 2004)*, Sept. 2004, Beijing, China.

Several papers in submission.

Unrefereed Technical Reports

- Dong Lu, Yi Qiao, Peter Dinda, and Fabian Bustamante, “Characterizing and Predicting TCP Throughput on the Wide Area Network”, Technical Report NWU-CS-04-34, Department of Computer Science, Northwestern University, April, 2004.
- Dong Lu, Yi Qiao, Peter Dinda, and Fabian Bustamante, “Modeling and Taming Parallel TCP on the Wide Area Network”, Technical Report NWU-CS-04-35, Department of Computer Science, Northwestern University, April, 2004.
- Yi Qiao, Dong Lu, Fabian Bustamante, and Peter Dinda, “Looking at the Server Side of Peer-to-Peer Systems”, Technical Report NWU-CS-04-37, Department of Computer Science, Northwestern University, March, 2004.
- Yi Qiao, and Fabian Bustamante, “The Effect of Lasting Friendships in P2P Protocols”, Technical Report NWU-CS-03-23 Department of Computer Science, Northwestern University, December, 2003.
- Yi Qiao, and Peter Dinda, “Network Traffic Analysis, Classification, and Prediction”, Technical Report NWU-CS-02-11, Department of Computer Science, Northwestern University, January, 2003.

Software

Peer Lifespan Traces: Over half-a-million peers' session length measurement and modeling traces. Available online at:
<http://www.aqualab.cs.northwestern.edu/lifeTrace.html>

TameParallelTCP: Estimate throughput and impact of parallel TCP flows. Released online at:
<http://plab.cs.northwestern.edu/Clairvoyance/Tame.html>

Recent Projects

MultiTorrents: MultiTorrents is a high performance P2P file sharing system that leverages incentivized protocol design to enhance system performance. MultiTorrents organizes the peers into a hierarchical tree structure according to peer capacity and encourages peers in the same layer to exchange data. To maintain strong incentives, MultiTorrents uses a partial *peer election* algorithm to build the tree structure and extends the "tit-for-tat" mechanism of BitTorrent to further encourage sharing. Our simulations show that MultiTorrents outperforms BitTorrent significantly.

Looking at the server side of P2P systems: In this work, we make the case for looking at the server side of peers, focusing on the problem of scheduling with the intent of minimizing the average response time experienced by users. We start by characterizing server workload based on extensive trace collection and analysis. We then evaluate the performance and fairness of different scheduling policies through trace-driven simulations. Our results show that average response time can be dramatically reduced by more effectively scheduling the requests on the server-side of P2P systems.

Lifespan-based organizational protocols and strategies for P2P systems: The high degree of peers' transiency seriously affect the scalability and efficiency of P2P systems. Through an independent measurement of over 500,000 peers' session lengths (i.e., lifespans), we found that peers' lifespan follows heavy-tailed pareto distribution. Based on this, we propose a number of illustrative lifespan-based overlay organizational protocols and query, caching and replication strategies that, by considering peers' lifespan as the first-key attribute, yielded P2P systems that are much more resilient to high degree of churn, with significantly improved query performance and system scalability. We demonstrated the effectiveness of our protocols and strategies through trace-driven simulation and wide-area experiments.

Multiscale predictability analysis of network traffic Network traffic prediction is crucial to many distributed application. The timescale of interest, however, is application-dependent, making it important to study how predictability of network traffic depends on the resolution of network traffic signal. In this work, we applied a wide range of linear and nonlinear time series models to a large number of packet traces, generating different resolution views of the traces through binning approach and wavelet-based approximations. We find that predictability is highly situational in practice and varies widely from trace to trace. Further, predictability does not always increase as the signal is smoothed, half of the time there is a "sweet spot", a particular resolution level where the traffic signal predictability is maximized. Our findings yield important implications for building effective adaptive distributed applications.

Characterizing and predicting TCP throughput on the Wide Area Network: Based on large-scale Internet-based measurement and experiments involving many sites distributed all over the world, we first characterized the transient End-to-End TCP throughput distribution, explored the strong correlation between TCP flow size and throughput, then based on the observations, we proposed a novel yet simple TCP throughput prediction model and a dynamic sampling rate adjustment algorithm. We evaluated the prediction system on the planet-lab.

Taming Parallel TCP on the Wide Area Network: Parallel TCP flows are broadly used in the high performance distributed computing community to enhance network throughput. There doesn't exist a practical model to predict the parallel TCP throughput and its effects on the cross traffic. To the best of our knowledge, we are the first to answer the following question on behalf of a user: *what number of parallel flows will give the highest throughput with less than a p% impact on cross traffic?* We evaluate our model through simulation-based and wide-area experimentation.

Size-based scheduling policies with inaccurate information: Size-based scheduling policies such as SRPT have been studied since 1960s and have been applied in various arenas. One important prerequisite to applying size-based scheduling is to know the sizes of all jobs in advance, which are unfortunately not always available. No work has been done to study the performance of size-based scheduling policies when only inaccurate scheduling information is available. We study the performance of SRPT and FSP as a function of the correlation coefficient between the actual job sizes and estimated job sizes.

Fat-tree based end-system multicast: FatNemo is a novel scalable peer-to-peer multi-source multicast protocol based on the idea of fat-trees. In fat-trees the available bandwidth increases as one moves up the tree, yielding a minimal mean and standard deviation of the response time. For many-to-many multicast applications, this eliminates the bottlenecks inside the overlay network. Our results show that FatNemo not only minimizes the average and the standard deviation of the response time, but also handles end host failures gracefully without suffering a performance penalty.

Service

Referee for IEEE MASCOTS, IEEE ICDCS.

Honors

19th Symposium on Operating Systems Principles (SOSP 2003) student travel grant, Northwestern University, Sept 2003.

Northwestern Student Travel Grant for LCR 2004 Conference, Northwestern University, Sept 2004.

National Aeronautics and Space Administration (NASA) ESS Fellowship, Northwestern University, 2002-2003.

Two-time Guanghua Scholarship, Tsinghua University, 1999-2001.

Nihon Kaiheiki Kogyo Co. (NKK) Fellowship, Tsinghua University, 1997-1998.

Three-time Outstanding Student Scholarship, Tsinghua University, 1993-1996.

References

Fabian E. Bustamante

Assistant Professor
Department of Computer Science
Northwestern University
1890 Maple Avenue
Evanston, IL 60201
fabianb@cs.northwestern.edu
(847) 491-2745 (voice)
(847) 491-5258 (fax)

Yan Chen

Assistant Professor
Department of Computer Science
Northwestern University
1890 Maple Avenue
Evanston, IL 60201
ychen@cs.northwestern.edu
(847) 491-4946 (voice)
(847) 491-5258 (fax)

Peter A. Dinda

Assistant Professor
Department of Computer Science
Northwestern University
1890 Maple Avenue
Evanston, IL 60201
pdinda@cs.northwestern.edu
(847) 467-7859 (voice)
(847) 491-5258 (fax)