

MICHAEL WALFISH

Department of Computer Science, The University of Texas at Austin
1 University Station C0500, Taylor Hall 2.124
Austin, TX 78712-0233

Fax: 512-232-7886
mwalfish@cs.utexas.edu
<http://www.cs.utexas.edu/~mwalfish>

EDUCATION

- 9/2002 – 11/2007 **Massachusetts Institute of Technology**, Cambridge, MA
Ph.D. in Computer Science, February 2008
S.M. in Computer Science, September 2004
Dissertation: “Defending Networked Resources Against Floods of Unwelcome Requests”
Advisor: Prof. Hari Balakrishnan
Minor: Probability and stochastic processes
- 9/1994 – 6/1998 **Harvard University**, Cambridge, MA
A.B. in Computer Science, *summa cum laude*, June 1998

RESEARCH INTERESTS

Networked computer systems. Sub-interests include security, network architecture, and performance.

EMPLOYMENT

- 1/2008 – Present **The University of Texas at Austin**, Austin, TX
Assistant Professor, 1/2009 – Present
Assistant Professor (on leave), 1/2008 – 12/2008
- 9/2008 – 12/2008 **University College London**, London, United Kingdom
Senior Research Associate
- 11/2007 – 8/2008 **Stanford University**, Stanford, CA
Postdoctoral scholar.
- Summer 2003,
Summer 2004 **International Computer Science Institute (ICSI)**, Berkeley, CA
Summer intern.
- 8/1999 – 5/2002 **Digital Fountain, Inc.**, Fremont, CA
Software engineer at computer networking startup.
- 9/1998 – 7/1999 **Oliver, Wyman & Company**, New York, NY
Consultant at financial services strategy consulting firm.
- Summer 1997 **Goldman Sachs & Company**, Boston, MA
Summer analyst in retail office, equities division.
- Summer 1996 **Microsoft Corp.**, Redmond, WA
Summer intern with Microsoft Word.

TEACHING

- Spring 2009 **Instructor** for UT CS course 395T, “Topics in Secure and Distributed Storage Systems”.
- Spring 2006 **Recitation Instructor** for MIT course 6.033, “Computer Systems Engineering”
Taught two sections of undergraduates. Each section met twice weekly. Recitations introduced new material and were a core part of this course. (The other recitation

instructors were faculty members.) Received 6.8/7.0 on anonymous student reviews, the highest rating of a staff member in this course in at least the past five years.

Fall 1996 **Teaching Assistant** for Harvard course CS50, “Introduction to Computer Science”
Received teaching award and 4.9/5.0 on anonymous student reviews.

Fall 1995 **Course Assistant** for Harvard course Math 1a, “Introduction to Calculus”

REFEREED CONFERENCE PUBLICATIONS

- M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenker. DDoS Defense by Offense. ACM SIGCOMM, September 2006.
- M. Walfish, J.D. Zamfirescu, H. Balakrishnan, D. Karger, and S. Shenker. Distributed Quota Enforcement for Spam Control. USENIX Symposium on Networked Systems Design and Implementation (NSDI), May 2006.
- M. Walfish, J. Stribling, M. Krohn, H. Balakrishnan, R. Morris, and S. Shenker. Middleboxes No Longer Considered Harmful. USENIX Symposium on Operating Systems Design and Implementation (OSDI), December 2004.
- H. Balakrishnan, K. Lakshminarayanan, S. Ratnasamy, S. Shenker, I. Stoica, and M. Walfish. A Layered Naming Architecture for the Internet. ACM SIGCOMM, September 2004.
- M. Walfish, H. Balakrishnan, and S. Shenker. Untangling the Web from DNS. USENIX Symposium on Networked Systems Design and Implementation (NSDI), March 2004.

REFEREED WORKSHOP PUBLICATIONS

- M. K. Aguilera and M. Walfish. No Time for Asynchrony. USENIX Workshop on Hot Topics in Operating Systems (HotOS), May 2009.
- M. Krohn, A. Yip, M. Brodsky, R. Morris, and M. Walfish. A World Wide Web Without Walls. ACM Workshop on Hot Topics in Networks (HotNets), November 2007.
- M. Walfish, H. Balakrishnan, D. Karger, and S. Shenker. DoS: Fighting Fire with Fire. ACM Workshop on Hot Topics in Networks (HotNets), November 2005.
- H. Balakrishnan, S. Shenker, and M. Walfish. Peering Peer-to-Peer Providers. 4th International Workshop on Peer-to-Peer Systems (IPTPS), February 2005.
- H. Balakrishnan, S. Shenker, and M. Walfish. Semantic-Free Referencing in Linked Distributed Systems. 2nd International Workshop on Peer-to-Peer Systems (IPTPS), February 2003.

TECHNICAL REPORTS AND WHITE PAPERS

- D. Mazières, S. Polu, and M. Walfish. Achieving the Potential of Provenance with Information Tracking. DARPA DOSR Workshop, July 2008.
- M. Walfish, J.D. Zamfirescu, H. Balakrishnan, D. Karger, and S. Shenker. Supplement to “Distributed Quota Enforcement for Spam Control”. Technical Report, MIT Computer Science and Artificial Intelligence Laboratory, MIT-CSAIL-TR-2006-033, April 2006.
- A. Gupta, M. Krohn, and M. Walfish. Can Basic ML Techniques Illuminate Rateless Erasure Codes? Technical Memo, MIT Laboratory for Computer Science, MIT-LCS-TM-643, May 2004.
- J. Considine, M. Walfish, and D. G. Andersen. A Pragmatic Approach to DHT Adoption. Technical Report, Boston University Computer Science Department, BUCS-TR-2003-024, December 2003.

THESES

- M. Walfish. Defending Networked Resources Against Floods of Unwelcome Requests. PhD Thesis, Department of Electrical Engineering and Computer Science, MIT, November 2007. Advisor: Prof. Hari Balakrishnan.
- M. Walfish. Semantic-Free Referencing in Linked Systems. Master's Thesis, Department of Electrical Engineering and Computer Science, MIT, September 2004. Advisor: Prof. Hari Balakrishnan.
- M. Walfish. An Analysis of the Metricom Radio Platform's Suitability for Interactive Audio Applications. Undergraduate Senior Thesis, Harvard College, April 1998. Advisor: Prof. H. T. Kung.

EXTERNAL TALKS

May 2009	“No Time for Asynchrony”, workshop talk at HotOS
November 2008	“Defending Networked Resources Against Floods of Unwelcome Requests”, invited talk at Max Planck Institute for Software Systems, Saarbrücken, Germany
November 2008	“Defending Networked Resources Against Floods of Unwelcome Requests”, invited talk at University of Cambridge, Cambridge, UK
July 2008	“Achieving the Potential of Provenance with Information Tracking”, DARPA workshop
September 2006	“DDoS Defense by Offense”, conference talk at SIGCOMM
September 2006	“DDoS Defense by Offense”, tech talk at Google
May 2006	“Distributed Quota Enforcement for Spam Control”, conference talk at NSDI
November 2005	“DoS: Fighting Fire With Fire”, workshop talk at HotNets
February 2005	“Peering Peer-to-Peer Providers”, workshop talk at IPTPS
December 2004	“Middleboxes No Longer Considered Harmful”, conference talk at OSDI
August 2004	“A Layered Naming Architecture”, workshop talk for IRTF HIP Research Group
March 2004	“Untangling the Web from DNS”, conference talk at NSDI

AWARDS

2003 – 2006	National Defense Science and Engineering Graduate (NDSEG) Fellowship
2003	National Science Foundation Graduate Research Fellowship
2002 – 2003	MIT Presidential Fellowship
1997	Phi Beta Kappa
1997	Certificate of Distinction in Teaching (for “Introduction to Computer Science” at Harvard)
1994 – 1998	John Harvard Scholarship (honorary)

PATENTS

2002	U.S. Patent Application No. 20020129159: M. Luby, R. Vainish, L. Rasmussen, D. Kushi, S. Simu, A. Perrig, R. Attias, M. Walfish, D. Hernek, J. Byers. <u>Multi-output packet server with independent streams</u> . (From work at Digital Fountain.)
2000	U.S. Patent No. 6,047,300: M. Walfish, D. Hachamovitch, R. Fein. <u>System and method for automatically correcting a misspelled word</u> . (From summer internship at Microsoft.)

SOFTWARE ARTIFACTS

(Names refer to projects listed on pages 5 and 6.)

- Speak-up project: the network element that sits in front of, and protects, the server. Release planned.
- DQE project: the distributed quota enforcer and client libraries. Released under GPL.
- DOA project: prototype of a DOA-enabled host and of two sample middleboxes that use DOA. (With two fellow students.)
- SFR project: prototype of a reference resolution infrastructure and of client libraries.

PROFESSIONAL SERVICE

- National Academies Panel on Digitization and Communications Science (2009–2011).
- Scribed NSDI 2007 PC meeting and wrote individual summaries of the PC discussions for the authors.
- Reviewer for NSDI (2007, 2008), SIGCOMM (2006, 2005), IPTPS (2006, 2005), SENSYS (2005).

PREVIOUS RESEARCH PROJECTS

COMBATING UNWANTED TRAFFIC

My dissertation covers two systems for combating unwanted traffic. *Speak-up* defends against application-level distributed denial-of-service (DDoS), and *Distributed Quota Enforcement (DQE)* controls spam volumes. For both of these unwanted traffic classes, the spurious traffic cannot be differentiated from the legitimate traffic reliably (*e.g.*, spam filters sometimes label legitimate email as spam and vice-versa). For this reason, speak-up and DQE do not *examine the contents* of requests (or email messages) but instead focus purely on *limiting the volume* of successful adversarial requests.

2005 – Present

Speak-up

Application-level DDoS overloads a server with legitimate-looking requests. With speak-up, a server so victimized encourages all clients to automatically send *higher* volumes of traffic. Speak-up supposes that attackers are already using most of their upload bandwidth and so cannot react to the encouragement, whereas good clients have spare upload bandwidth and so will send drastically higher volumes of traffic. As a result, the good clients will be much better represented in the traffic mix and will thus capture a much larger fraction of the server's resources than before. Experiments with a concrete instantiation of speak-up indicate that speak-up causes the protected server to spend resources on a group of clients in rough proportion to their aggregate upload bandwidth, which is the intended result. This work has appeared at HotNets 2005 and SIGCOMM 2006.

2005 – Present

Distributed Quota Enforcement (DQE)

Under DQE, each sender gets a quota of stamps and attaches a stamp to each email. Receivers communicate with a well-known *quota enforcer* to verify that the stamp on the email is fresh and to cancel the stamp to prevent reuse. The enforcer must be distributed, given its workload (between 100 and 200 billion emails are sent per day, which is between one and two million messages per second). The technical challenge is to design and implement the enforcer to handle this workload, to tolerate faults in its underlying machines, and to use as few machines as possible (to simplify management and costs). Our design for the enforcer requires only a few thousand machines and is notable for simplicity: mutually untrusting nodes implement a storage abstraction but avoid neighbor maintenance, replica maintenance, and heavyweight cryptography. This work has appeared at NSDI 2006.

NETWORK ARCHITECTURE

2003 – 2004

Delegation-Oriented Architecture (DOA)

There has long been a debate about middleboxes (intermediate network elements such as network address translators and firewalls) between purists, who decry them because they make the Internet less flexible, and pragmatists, who recognize that they are deployed for important reasons. DOA tries to satisfy both camps with an architectural extension to the Internet that retains middleboxes' useful functions while eliminating dangerous side-effects. The architecture involves two relatively modest changes to the status quo: (a) a set of references that are carried in packets and serve as persistent host identifiers, and (b) a way to resolve these identifiers to *delegates* chosen by the referenced host. A benefit of DOA is that middleboxes can be off-path, permitting functions such as off-path firewalls. This work appeared at SIGCOMM 2004 and OSDI 2004.

2002 – 2004

Semantic-Free Referencing (SFR)

Under SFR, Web objects are no longer linked to (*i.e.*, named) with DNS-based URLs but rather with semantic-free, human-unfriendly references. Because such links are location-independent and free of human semantics, these links (a) do not break when content changes administrative domains, (b) enable content to be seamlessly replicated across administrative domains, and (c) permit automatic management of the name space, in contrast to the status quo, in which a technical body (*i.e.*, ICANN) is in the business of arbitrating naming disputes. Web browsers resolve the references to network locations by contacting a resolution infrastructure based on distributed hash tables. This work appeared at IPTPS 2003 and NSDI 2004.