MICHAEL WALFISH

MIT Computer Science and AI Lab
The Stata Center, G982
32 Vassar Street
Cambridge, MA 02139

Phone: 617-233-7067 (cell), 617-253-7341 (office)
Fax: 617-253-8460
mwalfish@csail.mit.edu
http://nms.csail.mit.edu/~mwalfish

EDUCATION

September 2002 – Present
Massachusetts Institute of Technology, Cambridge, MA
Ph.D. candidate in Computer Science (degree expected summer 2007)
S.M. in Computer Science, September 2004
Dissertation topic: "Two Defenses Against 'Camouflage Attacks' "
Advisor: Prof. Hari Balakrishnan
Minor: Probability and stochastic processes

1994 – 1998
Harvard University, Cambridge, MA
A.B. in Computer Science, *summa cum laude*, June 1998
Phi Beta Kappa (top 4% of class)

RESEARCH INTERESTS

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

TEACHING EXPERIENCE

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"

EMPLOYMENT HISTORY

Summer 2003, Summer 2004
International Computer Science Institute (ICSI), Berkeley, CA
Summer intern. Worked on DOA and SFR research projects, described on pages 5 and 6.

August 1999 – May 2002
Digital Fountain, Inc., Fremont, CA
Software engineer at computer networking startup. The company's core technology is a proprietary erasure code. The company sold a server that was able to deliver 40,000 concurrent erasure-encoded UDP streams. End-hosts ran the Digital Fountain (DF) client to decode these streams and consume the content.

Responsible for design, implementation, and debugging of various software projects, including Linux device driver for smart NIC; PCI software running on the smart NIC; prototype of DF-over-TCP and resulting comprehensive changes to Linux TCP stack; media-on-demand client; and automatic installation system for DF's client software.

September 1998 – July 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; worked in equities division.

Summer 1996 **Microsoft Corp**., Redmond, WA
Summer intern with Microsoft Word. Researched, designed, and prototyped an original Word 2000 feature: automatic correction of misspellings not in the AutoCorrect list.

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. 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

- 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. 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.

TALKS

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	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. Multi-output packet server with independent streams. (From work at Digital Fountain.)
2000	U.S. Patent No. 6,047,300: M. Walfish, D. Hachamovitch, R. Fein. System and method for automatically correcting a misspelled word. (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

- Scribed NSDI 2007 PC meeting and wrote individual summaries of the PC discussions for the authors.
- Reviewer for NSDI (2007), SIGCOMM (2006, 2005), IPTPS (2006, 2005), SENSYS (2005).

ACTIVITIES AND PERSONAL INTERESTS

Co-captain of CS lab's softball team, distance running (ran 2006 Cape Cod Marathon), cooking, and literature.

REFERENCES

Prof. Hari Balakrishnan
MIT Computer Science & AI Lab
32 Vassar Street, 32G-940
Cambridge, MA 02139
(617) 253-8713
hari@csail.mit.edu

Prof. Scott Shenker
UC Berkeley
465 Soda Hall
Berkeley, CA 94720
(510) 643-3043
shenker@icsi.berkeley.edu

Prof. M. Frans Kaashoek
MIT Computer Science & AI Lab
32 Vassar Street, 32G-992
Cambridge, MA 02139
(617) 253-7149
kaashoek@csail.mit.edu

Prof. David Karger
MIT Computer Science & AI Lab
32 Vassar Street, 32G-592
Cambridge, MA 02139
(617) 253-6167
karger@mit.edu

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.