

Michael F. Schwartz

Phone: (206) 683-0392

Email: mfschwartz@gmail.com

Web: www.CodeOnTheRoad.com

Vitae Last Updated: October 5, 2009

Professional Experience

August 2008-present: Member of Technical Staff, Google, Inc.

June 2007-July 2008: President, Code On The Road, LLC.

Independent consultant providing technical research and development services, with particular emphasis on large scale search, content extraction, and other back-end service infrastructure problems.

June 2004-May 2007: Vice President of Technology Innovation for Local Matters, Inc.

Implemented search services for national Yellow Pages sites; built out operational support infrastructure for deploying FAST search technology; co-directed Local Matters labs, performing performance analysis and research and development of crawling, content extraction, ontology-based search, and yellow pages/white pages service integration.

March 2000-May 2004: President, Code On The Road, LLC.

Independent consultant providing technical research and development and legal expert witness services on problems related to scalable network and service infrastructure. Projects included work on network and data management architecture, caching and data replication, search, security, transport protocols, network instrumentation and measurement, and performance analysis.

November 1998-March 2000: Vice President of Engineering and Chief Technology Officer, Affinia, Inc.

Co-founded and helped raise seed funding for the company, built the engineering team, led development on key parts of the technical infrastructure, and brought company's initial database-driven e-commerce service to market.

October 1997 to November 1998: Director of Server Engineering and Senior Scientist, @Home Network.

Managed 25 person team responsible for global @Home server architecture, server product development, core services (mail, news, DHCP, Radius, etc.), server hardware and operating systems, capacity planning, data distribution and caching, subscriber and modem provisioning, usage data analysis, NT platform integration, and server and network management.

November 1995 to October 1997: Senior Scientist, @Home Network.

Led architecture and deployment of distributed subscriber provisioning system. Implemented initial search and directory services. Co-architected media partner content replication service. Performed packet trace- and cache-log based global network architecture assessment, spearheading effort to enhance performance and effectiveness of the network and server infrastructure. Produced network traffic models used for capacity planning. Implemented cable throughput and packet loss test system.

April 1994 to October 1995: Associate Professor, University of Colorado, Boulder.

Led design and managed implementation, integration, and deployment of the *Harvest* search and caching system. Compared with previous systems, the *Harvest* indexing system reduced CPU load by a factor of 6,000; network traffic by a factor of 60; and disk usage by a factor of 40. The *Harvest* cache was commercialized and acquired by Network Appliance, and also formed the basis for the freeware *Squid* cache. The *Harvest* indexing system has been commercialized by Netscape in their *Compass* product. Two *Harvest* protocols (SOIF and ICP) underwent standardization efforts in the Internet Engineering Task Force.

August 1987-April 1994: Assistant Professor, University of Colorado, Boulder.

Prior to above *Harvest* project, led design and implementation of several other Internet information systems, including *Netfind* (a user directory service), *Fremont* (a network management service), and *Essence* (a customizable content extraction system later incorporated into *Harvest*). Also carried out several measurement studies regarding size, traffic patterns, and uses of the Internet.

October 1985-August 1987: Research Assistant, University of Washington.

Implemented heterogeneous name service and remote procedure call subsystems in the *Heterogeneous Computer Systems* Project. Also implemented modifications to Berkeley's implementation of the Domain Name Service (BIND) to support dynamic updates and data of unspecified type.

June-October 1985: Member of the Technical Staff, Bell Communications Research.

Implemented testbed environment for experimenting with multi-vendor telephony services.

September 1983-June 1985: Research Assistant, University of Washington.

Ported interprocess communication subsystem for the *Eden* distributed system from home-grown system to recently released BSD 4.1 sockets. Implemented and measured high performance file transfer system, and reliable broadcast system.

June 1980-September 1983: Member of the Technical Staff, Hughes Aircraft Company.

Implemented radar system debugging and test environment. Implemented radar targeting software testbed. Led design of real-time radar data processing operating system.

Selected Professional Positions

Technical Advisory Board member, Peribit Technologies, 2001-2005.

Technical Advisory Board member, Internet Archive Project, August 1999-May 2000.

Internet Advisory Board member, Vital Signs Software, Inc., January 1998-January 2000.

Program committee member, USENIX '98 Conference.

Contributed to technical development proposal presented by Netscape's Public Policy Counsel to the World Intellectual Property Organization at December 1996 Geneva meeting concerning Internet copyright issues.

Co-chair, World Wide Web Consortium Distributed Indexing/Searching Workshop, Cambridge, MA, May 1996.

Program committee member, SIGCOMM '95 and '96 Conferences.

Program committee member, HPCC/IITA Digital Library Workshop, McLean, VA, May 1995.

Chair, U.S. InterNIC midterm review panel, Arlington, VA, November 1994.

Panelist, ARPA/NASA/NSF Digital Library Initiative review panels, April 1994.

Editor, IEEE/ACM Transactions on Networking, 1992-1996.

Guest Editor, Special Issue on the Global Internet, IEEE Journal of Selected Areas in Communication, 1993-1995.

Member, Internet Society Editorial Advisory Board, 1991-1996.

Chair, Internet Research Task Force Research Group on Resource Discovery (1991-1996). Also, member of the Internet Survey Working Group (1995), and of the Internet Research Steering Group (1991-1996).

Program committee member, Second and Fourth International World Wide Web Conferences, 1994-1995.

Reviewer for U.S. Congressional Office of Technology Assessment, U.S. National Science Foundation, U.S. National Telecommunications and Information Administration, and Canadian National Research Council, 1988-1995.

Selected Research Grants and Contracts

Co-Principal Investigator, NSF: "A Distributed Testbed for National Information Provisioning", joint with the San Diego Supercomputer Center, \$390,000, 1995-1998. (Left project in November 1995.)

Co-Principal Investigator, NSF: "High Performance Infrastructure for Computational Science", \$1,900,000, 1995-1999.

Principal Investigator, ARPA: "Agile Integration of Multiprotocol, Multimedia Information Services", \$2,100,000, with subcontracts to the Internet Multicasting Service and the Oregon Graduate Institute of Science & Technology, 1995-1998.

Principal Investigator, ARPA: "New Internet Services for Scalable Resource Discovery", \$1,800,000, with subcontracts to Transarc Corp., the University of Southern California, and the University of Arizona, 1993-1996.

Principal Investigator, NSF: "A Measurement Study of Changes in Service-Level Reachability in the Global TCP/IP Internet", \$130,000, 1992-1995.

Selected Talks and Panels

Invited talk, Stanford University, Palo Alto, CA, April 1998.

Invited panelist, InfoCom conference, San Francisco, CA, March 1998.

Keynote talk, Multimedia Computing and Networking conference, San Jose, CA, January 1998.

Invited talk, Workshop on Frontiers in Distributed Information Systems, Aspen, CO, June 1997.

Invited talks, Xerox Palo Alto Research Center (Palo Alto, CA), University of Washington (Seattle, WA), University of Colorado (Boulder, CO), April-May 1997.

Invited panelist, Network/Interop Conference, Las Vegas, NV, May 1997.

Invited talk, Intel Emerging Communications Forum, Portland, OR, August 1996.

Invited talk, National Research Council Workshop on Materials and Processes Research and the Information Highway, Palo Alto, CA, April 1996.

Invited talks, Berkeley and Stanford digital library seminars, Berkeley and Palo Alto, CA, March 1996.

Invited talk, Netscape Developer Conference, San Francisco, CA, March 1996.

Selected Publications

1. E. Dixon, H. Franklin, J. Kint, G. Klyne, D. New, S. Pead, M. Rose, M. Schwartz, The Application Exchange (APEX) Option Party Pack, Part Deux!, RFC 3342, July 2002.
2. M. Schwartz, The ANTACID Replication Service: Rationale and Architecture, draft-schwartz-antacid-service-00 (work-in-progress IETF standards-track specification), October 2001.
3. M. Schwartz, M. Rose, K. Carlberg, The APEX Publish-Subscribe Service, draft-schwartz-apex-pubsub-01 (work-in-progress IETF standards-track specification), October 2001.
4. D. R. Hardy and M. F. Schwartz. Customized Information Extraction as a Basis for Resource Discovery. *ACM Transactions on Computer Systems*, 14(2), pp. 171-199, 1996.
5. A. Chankhunthod, P. B. Danzig, C. Neerdaels, M. F. Schwartz, and K. J. Worrell. A Hierarchical Internet Object Cache. Proceedings of the USENIX Technical Conference, pp. 153-163, San Diego, California, January 1996.
6. C. Mic Bowman, Peter B. Danzig, Darren R. Hardy, Udi Manber and Michael F. Schwartz. The Harvest Information Discovery and Access System. *Computer Networks and ISDN Systems*, 28, pp. 119-125, 1995. (Slightly extended version of an earlier conference paper, selected as among the best papers at the Oct '94 *International World Wide Web Conference*.)
7. J. D. Guyton and M. F. Schwartz. Locating Nearby Copies of Replicated Internet Servers. *Proceedings of the SIGCOMM Conference*, pp. 288-298, Cambridge, MA, August 1995.
8. M. F. Schwartz and C. Pu. Applying an Information Gathering Architecture to Netfind: A White Pages Tool for a Changing and Growing Internet. *IEEE/ACM Transactions on Networking*, 2(5), pp. 426-439, October 1994.
9. B. Chhabra, D. R. Hardy, A. Hundhausen, D. Merkel, J. Noble and M. F. Schwartz. Integrating Complex Data Access Methods into the Mosaic/WWW Environment. *Proceedings of the Second International World Wide Web Conference*, pp. 909-919, Chicago, IL, October 1994.
10. C. M. Bowman, P. B. Danzig, U. Manber and M. F. Schwartz. Scalable Internet Resource Discovery: Research Problems and Approaches. *Communications of the ACM*, 37(8), pp. 98-114, August 1994.
11. J. D. Guyton and M. F. Schwartz. Experiences with a Survey Tool for Discovering Network Time Protocol Servers. *Proceedings of the USENIX Summer Conference*, pp. 257-265, Boston, Massachusetts, June 1994.
12. P. B. Danzig, R. S. Hall and M. F. Schwartz. A Case for Caching File Objects Inside Internetworks. *Proceedings of the SIGCOMM Conference*, pp. 239-248, San Francisco, California, September 1993.
13. M. F. Schwartz and J. S. Quarterman. The Changing Global Internet Service Infrastructure. *Internet Research: Electronic Networking Applications and Policy*, 3(3), pp. 8-25, Fall 1993.
14. M. F. Schwartz and D. C. M. Wood. Discovering Shared Interests Using Graph Analysis. *Communications of the ACM*, 36(8), pp. 78-89, August 1993.
15. C. M. Bowman, P. B. Danzig and M. F. Schwartz. Research Problems for Scalable Internet Resource Discovery. *Proceedings of INET '93*, pp. DFB1-DFB10, San Francisco, California, August 1993.
16. M. F. Schwartz and J. S. Quarterman. A Measurement Study of Changes in Service-Level Reachability in the Global Internet. *Proceedings of INET '93*, pp. CBB1-CBB10, San Francisco, California, August 1993.
17. D. C. M. Wood, S. S. Coleman and M. F. Schwartz. Fremont: A System for Discovering Network Characteristics and Problems. *Proceedings of the USENIX Winter Conference*, pp. 335-348, San Diego, California, January 1993.

Education

Ph.D., Computer Science, University of Washington, 1987.

M.S., Computer Science, University of Washington, 1985.

B.S. *cum laude*, Mathematics-Computer Science, UCLA, 1982.