

Arup Mukherjee

San Jose, CA 95132; (408) 347-1354

<http://www.mukherjee.net/arup/>
arup@mukherjee.net

Professional Interests

Building high-performance, scalable systems in a startup or advanced development environment. I am interested in distributed systems, operating systems, electronic commerce, networked applications, mobile computing, and computer security.

Education

Carnegie Mellon University: Doctor of Philosophy (**Ph.D.**) degree in Computer Science.

November 1998

Thesis work: Developed techniques and infrastructure for building online services in the face of severe bandwidth constraints. Designed, implemented and evaluated the Oasis system, wherein Web-based Java agents and services hosted by a network of programmable proxy servers cooperate to optimize communication by sharing data and capabilities within the network and at end-user devices. Oasis implementation was targeted at services deployed on the World Wide Web for access via dialup and mobile clients. Thesis advisor: Professor Daniel P. Siewiorek.

Carnegie Mellon University: Master of Science (**M.Sc.**) degree in Computer Science.

May 1993

Coursework: programming languages, operating systems, computer systems, algorithms, parallel algorithms, distributed programming environments, and artificial intelligence. Noteworthy seminar courses attended: Mach 3.0 kernel internals, mobile computing.

Research: Developed and implemented a means of benchmarking the robustness of complex software systems. (Paper in IEEE Transactions on Software Engineering, June 1997.) Conducted as part of the Fault Tolerant Mach effort, with Professor Daniel P. Siewiorek.

University of Pennsylvania: Bachelor of Science in Engineering (**B.S.E.**) degree in Computer Science & Engineering. *May 1990*
Graduated **Summa Cum Laude**, with additional General Honors distinction.

Senior Project: Design and prototyping of the hardware & software for a MC68008/X-Windows based voicemail system for workstations.

Coursework: Computer science and engineering curriculum tailored for strong emphasis on software and hardware systems.

GPA: 4.00/4.00 (*Computer Science & Engineering*), 3.82/4.00 (*Cumulative*)

Experience

Yahoo! Shopping Engineering, Yahoo!, Inc, Santa Clara, CA.

May 1999-present

Senior Software Engineer, Remote Merchant Integration group: One of a team of four engineers who designed and implemented the programmable proxy server infrastructure used to transparently integrate the web sites of external clients into Yahoo! Shopping. The proxy system adds Yahoo! branding, integration with Yahoo! services (such as Yahoo! wallet), and extracts and logs all transactions on the fly while servicing several million daily pageviews to over 200 large brand-name merchants. Was responsible for several large subsystems written from scratch in C++, and was also heavily involved in operationally launching the system and in developing all required deployment, monitoring, management and server redundancy mechanisms. Promoted to technical lead and manager of this team in February, 2000 and to senior manager in October, 2001.

Technical Lead and Manager / Senior Manager, Remote Merchant Integration group: Led efforts to greatly increase the stability, scalability, flexibility and efficiency of the Remote Merchant Integration proxy system while adding support for new merchants and new services, including several applications from outside shopping. Duties included architecting many system changes to attain the aforementioned goals, guiding and mentoring up to 8 other engineers, hands-on development of technically complex parts of the system, management of the operational servers, leading the internationalization effort, supporting 9 international deployments and participating in recruiting efforts.

Computer Science Department, Carnegie Mellon University, Pittsburgh, PA.

Aug 1991-March 1999

Research Fellow, Teaching Assistant, and Volunteer Systems Programmer and Administrator: Designed, implemented, and evaluated a system to benchmark the fault tolerance of operating systems platforms. Subsequently, designed and implemented an infrastructure of enhanced Web proxies for the efficient deployment of Java online services in bandwidth-constrained environments. Served as a teaching assistant for courses in operating systems, graduate distributed systems, and graduate computer systems. *Selected voluntary efforts:* ported the in-house remote filesystem backup mechanism to Linux and AIX systems; installed and supported XEmacs locally for a user base of over 100 users using over 10 different workstation environments; was instrumental in the installation, porting, and support of a large application software base for a local network of about 30 AIX 3.2 workstations.

Spring Project, Sun Microsystems Labs, Mountain View, CA.

Summer 1993

Summer Research Intern: Converted the distributed filesystem of SpringOS (an experimental object-oriented microkernel operating system) to a stackable filesystem composed of file encapsulation and coherency layers. Enhanced the coherency support to handle arbitrary stacked filesystem configurations. Was awarded a vice president's letter of recognition and salary bonus for my work.

Spring Project, Sun Microsystems Labs, Mountain View, CA.

Summer 1992

Summer Research Intern: Analyzed the virtual memory system of SpringOS. Characterized its performance, added functionality and significant performance improvements. Was awarded a vice president's letter of recognition and salary bonus for my accomplishments.

Large Systems, IBM Research, Thomas J. Watson Research Center, Hawthorne, NY.

May 1990 - September 1991

Systems Programmer/Researcher: Investigated the applicability of affinity-based scheduling to Unix-based parallel processing (paper in Winter Usenix, 1992), and thrashing in real-addressed caches due to virtual memory management. Also implemented frequency based buffer-cache

management in the 4.3 BSD filesystem. *Parallel AIX Fileserver Project*: Designed and implemented much of the low-cost disk sharing scheme in the prototype via kernel extensions, device drivers, and modifications to the journalled filesystem under AIX 3.1 on the IBM RS/6000.

General Robotics and Active Sensory Perception Laboratory, University of Pennsylvania, Phila., PA. *June 1987 - May 1990*

Systems Programmer: Wrote and maintained a large extension to the device independent image processing software library in use at the lab. Work included network support using TCP/IP, and design and implementation of a set of device independent output routines for X windows (X10R4 thru X11R4), POSTSCRIPT printers, and IKONAS and DATA TRANSLATION frame buffers.

Highly Parallel Systems, IBM Research, Thomas J. Watson Research Center, Hawthorne, NY. *Summer 1989*

Systems Programmer and Administrator: Maintained a network of IBM RTs running Mach. Modified the Mach kernel to meet local needs, assembled and installed new equipment, and installed and supported the Andrew File System and the X Window System.

Teaching Experience

Carnegie Mellon University

- Teaching assistant, undergraduate *Operating Systems*. Course project assignment was the implementation of an operating system, including virtual memory and file system, for a SPARC-like virtual machine. Awarded the School of Computer Science graduate student teaching award for my work in this course. *Fall 1995*
- Teaching assistant, graduate *Distributed Systems*. Course project assignment was the design and implementation of a system for conducting secure electronic commerce over the World Wide Web. *Spring 1994*
- Teaching assistant, graduate *Computer Systems*. Responsibilities included grading and homework design. *Fall 1992*

Honors and Awards

Carnegie Mellon University School of Computer Science Graduate Student Teaching Award for "excellence in teaching and dedication to students and to the academic community." *August 1996*

Letters of recognition and salary bonuses for contributions to the Spring project at Sun Microsystems Labs. *August 1992, August 1993*

Graduated Summa Cum Laude, with General Honors distinction for completing 8 honors courses. *May 1990*

Honors Society Memberships: Sigma Xi, Phi Beta Kappa, Eta Kappa Nu, Tau Beta Pi. *respectively 1991, 1990, 1989, 1988 - present*

Benjamin Franklin Scholar at the University of Pennsylvania (special honors program for top 7%.) *January 1988 - May 1990*

Dean's List, University of Pennsylvania. *1988, 1989, 1990*

Second place, Class of 1880 mathematics exam, open to all freshmen at the University of Pennsylvania. *May 1987*

Graduation award for "setting academic standards for others to strive to emulate" from the International School of Lusaka. *June 1986*

First prize, annual National Mathematical Problems Contest, sponsored by the Zambia Association for Mathematics Education. Only person ever in the history of the contest to win twice. *1985 and 1986*

Publications

A. Mukherjee, "Supporting Online Services in Environments Constrained by Communication," *Ph.D. Thesis, CMU CS Tech Report CMU-CS-98-172*, November 1998.

A. Mukherjee and D. P. Siewiorek, "Measuring Software Dependability by Robustness Benchmarking," *IEEE Transactions on Software Engineering*, June 1997. (Preliminary version, *CMU CS Tech Report CMU-CS-94-148*, was published in 1994.)

M. Devarakonda, A. Mukherjee and W. Kish, "Meta-Scripts as a Mechanism for Complex Web Services," *Proceedings of the Fifth Workshop on Hot Topics in Operating Systems*, May 1995.

A. Mukherjee and D. P. Siewiorek, "Mobility: A Medium for Computation, Communication and Control," *IEEE Workshop on Mobile Computing Systems and Applications*, December 1994.

M. Devarakonda and A. Mukherjee, "Issues in Implementing Cache-Affinity Scheduling," *Proceedings of the Winter Usenix 1992*, January 1992.

Computer Skills

Languages: C++, C, Java, Perl 5, HTML, Tcl/Tk, Pascal, shell scripts (csh and sh), Lisp, Scheme.

WWW: Java, HTML, HTTP, Apache, Javascript, proxy server design, server farm redundancy and load balancing, caching, some DHTML.

Operating Systems: Thorough knowledge of Unix systems including Linux, FreeBSD 2.2 and 4.x, AIX 3.2, BSD 4.3, SunOS, Solaris, and Mach. Some knowledge of Windows platforms.

Unix Kernel: I have worked with the internals of AIX 3.1, BSD 4.3, and Mach 2.5 in the areas of scheduling, filesystem (including the AIX journalled file system) and virtual memory management. I have modified kernels directly in these areas, and have also written device drivers, virtual file systems, and system calls.

Unix Software: most standard unix tools, including apache, perl, X11, NFS, kerberos 4, T_EX, L_AT_EX, TCP/IP, bash/csh/sh, sed

Personal

- I am a Canadian citizen, a native speaker of English, and a legal permanent resident of the United States.
- I studied French for several years, and retain limited ability in spoken and written French. I also speak Bengali fluently.