

Varun Gupta

Computer Science Department
4212 Wean Hall
Carnegie Mellon University
Pittsburgh, PA 15213
USA

Phone: 412.268.1188
Fax: 412.268.5576
varun@cs.cmu.edu
<http://www.cs.cmu.edu/~varun>

Research Interests

Performance modeling, design and analysis of computer systems
Algorithms for distributed systems/computer networks
Stochastic modeling

Education

Ph.D. Computer Science, Expected 2010
Carnegie Mellon University, Pittsburgh, PA, USA
Advisor: Prof. Mor Harchol-Balter

B.Tech. Computer Science and Engineering, *President's Gold Medal winner*, August 2004
Indian Institute of Technology, New Delhi, India
Thesis: *Algorithms for Computational Biology: Sequence Analysis*
Advisor: Prof. S.N. Maheshwari

Research Projects and Experience

Carnegie Mellon University
Graduate Assistant
Prof. Mor Harchol-Balter

Pittsburgh, PA
August 2004–Present

My research interests lie in the design of resource allocation and scheduling schemes for performance-oriented operation of server farms and distributed/P2P systems. Currently, my research involves stochastic modeling and analysis of queueing networks and scheduling policies to obtain an intuitive understanding of the effects of the system parameters on various performance metrics. A brief description of a few selected problems I have worked on follows.

Load-balancing heuristics: This project involves improving the performance of common load-balancing server farm architectures, comprising of a high-speed front-end router/-dispatcher (such as the IBM Network Dispatcher). Here we analyze different routing policies for the router and show the near-optimality of a simple Join-the-Shortest-Queue heuristic with Processor sharing front end servers. This work appeared in Performance 2007.

Resource sharing systems: In this project, we consider a time sharing (processor sharing) computer system where the server splits its capacity over the jobs in its buffer. This model is applicable in all kinds of computing centers, from web servers and database servers to communication networks. However, in reality there is always a limit on how many jobs can share the processor at one time. Here, the goal is to analyze the effect of the degree of sharing on the performance of the system and possibly come up with better sharing mechanisms.

Systems with time-varying load: One of my earlier project involves analyzing computer system performance under situations where the workload and service capacity vary

over time. Understanding of systems with time varying traffic and service characteristics turns out to be a fundamental building block for analyzing more complicated resource sharing systems (e.g. cycle stealing) and multi-server systems. This research involves inventing new queueing-theoretic techniques and approximations to handle time-varying load, and appeared in ACM SIGMETRICS 2006.

Bell Laboratories, Alcatel-Lucent
Research Intern
Host: Sem Borst

Murray Hill, NJ
Summer 2008

The focus of the internship was to devise simple distributed algorithms for managing a hierarchical content distribution network to use the cache at each node in the hierarchy most efficiently while minimizing the access latencies. In this work, we characterize the structure of the optimal centralized caching scheme, propose simple distributed algorithms with provable performance guarantees compared to the optimal centralized scheme, and perform stochastic analysis of some of the proposed distributed algorithms.

Microsoft Research Ltd.
Research Intern
Host: Milan Vojnovic

Cambridge, UK
Summer 2006

The main goal of this internship was to investigate best sampling strategies for epidemic-style information dissemination to improve effectiveness of computer patching. I also investigated alternate architectures for Windows Update - one of the largest content distribution networks in the world. Part of this work will appear in IEEE INFOCOM 2008.

Indian Institute of Technology
Bachelor's Thesis in Bioinformatics
Prof. S.N. Maheshwari

New Delhi, India
Aug 2003–Jul 2004

For my B.Tech. thesis, I worked under Prof. Maheshwari's supervision in collaboration with Prof. Alok Bhattacharya (School of Life Sciences, Jawaharlal Nehru University) to identify and solve problems in the field of Computational Biology. We developed a heuristic algorithm to find approximate repeats in DNA sequences. We also designed a Fourier-transform based algorithm to find approximate occurrences of a pattern (in terms of *Hamming distance*) in a database string with a general alphabet distance/similarity matrix.

Report available at <http://www.cs.cmu.edu/~varun/docs/report.pdf>

Read-Ink Tech. Pvt. Ltd.
Handwriting Recognition

Bangalore, India
Summer 2003

Prof. Thomas O. Binford (Emeritus Professor of Computer Science - Research, Stanford University)

During my summer internship, I worked with Prof. Binford and his team who are developing the next generation handwriting recognition engine with a target recognition rate of 99%.

Publications

- Fundamental Characteristics of Queues with Fluctuating Load with M. Harchol-Balter, A. Scheller-Wolf, U. Yechiali, *ACM SIGMETRICS/Performance 2006*.
- The Effect of Higher Moments of Job Size Distribution on the Performance of an $M/G/s$

Queueing System

with J. Dai, M. Harchol-Balter, B. Zwart, *Performance Evaluation Review*, September 2007.

- Analysis of Join-the-Shortest-Queue Routing for Web Server Farms with M. Harchol-Balter, K. Sigman, W. Whitt, *Performance* 2007.
- Sampling Strategies for Epidemic-style Information Dissemination with M. Vojnovic, T. Karagiannis, C. Gkantsidis, *IEEE INFOCOM* 2008.
- Finding the Optimal Quantum Size: Sensitivity Analysis of the $M/G/1$ Round-Robin Queue To appear in *Performance Evaluation Review*, 2008.
- Fluid Level in a Reservoir with an On-Off Source with P. Harrison, To appear in *Performance Evaluation Review*, 2008.

Teaching

Teaching Assistant

Fall 2006

15-451: Algorithm Design and Analysis. (Undergraduate Course)

Instructor: Prof. Avrim Blum and Prof. Manuel Blum

Teaching Assistant

Spring 2005

15-849: Performance Modeling. (Graduate Course)

Instructor: Prof. Mor Harchol-Balter

Awards and Honors

- Awarded the President's Gold Medal for the highest CGPA at the end of undergraduate studies, Indian Institute of Technology, New Delhi, 2004.
- Awarded the R. Vibhakar Memorial Award for the best student of third year during the 2002-03 session, Indian Institute of Technology, Delhi.
- Awarded the Rajiv Bhambawale Trust Award for securing highest CGPA amongst the third year undergraduate students during 2002-03 session, Indian Institute of Technology, Delhi.
- Awarded the Raman Subramanian Award for securing the highest CGPA amongst the second year undergraduates during the 2001-02 session, Indian Institute of Technology, Delhi.
- Nominated as an executive member of the Delhi wing of IEE (UK) for the year 2001.
- Secured All India Rank 22 at the Indian Institute of Technology Joint Entrance Examination (IITJEE 2000) and All India Rank 10 at Roorkee Entrance Test 2000 (now IIT Roorkee).
- Secured 2nd rank (senior level) in Delhi region at the Regional Mathematics Olympiad 1997 organised by National Board for Higher Mathematics (NBHM). Consequently, I was selected for the Nurture Program to receive training for Indian National Mathematics Olympiad (INMO).
- Secured 17th rank (junior level) in Delhi region at Regional Mathematics Olympiad 1996 organised by National Board for Higher Mathematics (NBHM).

Skills

- Languages: English, Hindi, basic French
- Computer Skills: C/C++, Java, SML, Unix, Linux, L^AT_EX, Matlab, Mathematica, Perl, HTML, SQL