GURMEET SINGH MANKU

530 Lawrence Expy #521, Sunnyvale, CA 94085 (650) 810 6863
manku@cs.stanford.edu http://www.cs.stanford.edu/~manku

RESEARCH INTERESTS

Algorithms for massive datasets — randomized algorithms — database systems — peer-to-peer systems.

EDUCATION

Stanford University	Ph. D. in Computer Science	2004	GPA 4.00
University of California at Berkeley	M. S. in Computer Science	1997	GPA 3.97
Indian Institute of Technology Delhi	B. Tech. in Computer Science	1995	GPA 9.91

EMPLOYMENT

Google Inc	Software Engineer	Aug 2004 – present
Gigabeat Inc (acquired by Napster)	Summer Intern	Summer 2000
IBM Almaden Research Center	Staff Software Engineer	July 1997 – Sep 1999
Intel Development Labs (Strategic CAD Technology)	Summer Intern	Summer 1996

HONORS AND AWARDS

Stanford Graduate Fellowship, Stanford University, 1999-2003.

University Fellowship, U C Berkeley, 1997.
U C Regents Fellowship, U C Berkeley, 1995-96.

ACM International Collegiate Programming Contest, 1996-97, UC Berkeley team member.
Annual Berkeley Programming Contest, 1996, Third rank.

3rd among 350+ students of all disciplines, 2nd among 45 students in Computer Science, IIT Delhi, 1995.
Award for the Best Undergraduate Project in Computer Science, IIT Delhi, 1995.
R Vibhakar Cash Award for best student in all disciplines at IIT Delhi, 1993-94.
R Bambawale Prize and R Subramanian Award for best student in all disciplines at IIT Delhi, 1992-93.

Indian National Mathematics Olympiad - 1990, among the top 20 students in India.

PATENTS

W J LABIO, G T NGUYEN, W W LIU, G S MANKU — System and Method for Searching Peer-to-Peer Computer Networks by Selecting a Computer Based on At Least a Number of Files Shared by the Computer — US Patent #07089301, Issued: Aug 8, 2006.

B G LINDSAY, G S MANKU, S RAJAGOPALAN — Single Pass Space Efficient System and Method for Generating an Approximate Quantile in a Data Set Having an Unknown Size — US Patent #06343288, Issued: Jan 29, 2002

B G LINDSAY, G S MANKU, S RAJAGOPALAN — Single Pass Space Efficient System and Method for Generating Approximate Quantiles Satisfying an a priori User-Defined Approximation Error — US Patent #06108658, Issued: Aug 22, 2000

MISCELLANEOUS ACTIVITIES

Program Committee Member, VLDB 2007.
Invited chapter in “Data-Stream Management”, a book edited by Garofalakis, Gehrke and Rastogi, 2004.
External reviewer for SIGMOD, VLDB, PODS, ICDE, UbiComp and USENIX, 1999 – 2004.
Organizer of the Stanford/ACM Annual Programming Contest – Oct 2003.

TEACHING

Teaching assistant for CS-145 (Database Systems), Stanford University, Fall 2002.
Guest lecturer for CS-361A (Advanced Algorithms and Data Structures), Stanford University, Spring 2002.
Teaching assistant for CS-154 (Automata and Complexity Theory), Stanford University, Winter 2003.

TALKS

ID Management and Load Balance in Distributed Hash Tables – PODC 2004 – 27 July, 2004.
Distributed Hash Tables: A Modular Approach (Ph.D. Defense) – Stanford University – 2 June, 2004.
Distributed Hash Tables: A Modular Approach (Poster) – SNRC Workshop 2004 – 27 May, 2004.
Optimal Routing in Chord – SODA 2004 – New Orleans, LA – 11 Jan, 2004.
Randomized ID Selection and Routing in DHTs – P2P Research Group at U C Berkeley – 17 Nov, 2003.
SETS: Search Enhanced by Topic Segmentation – SIGIR 2003 – Toronto, Canada – 31 Jul, 2003.
Routing Networks for Distributed Hash Tables – PODC 2003 – Boston, MA – 14 Jul, 2003.
Symphony: Distributed Hashing in a Small World – USITS 2003 – Seattle, WA – 27 Mar, 2003.
Symphony: A DHT using Small-World Routing – P2P Research Group at U C Berkeley – 10 Mar, 2003.
Histograms and Quantiles – Guest Lecture in course CS-361A at Stanford University – 9 Oct, 2002.
Approximate Frequency Counts over Data Streams – VLDB 2002 – Kowloon, Hongkong – 21 Aug, 2002.
Random Sampling Techniques for Approximate Quantiles – SIGMOD 1999 – Philadelphia, PA – Jun 1999.
Approximate Quantiles in One Pass and with Limited Memory – SIGMOD 1998 – Seattle, WA – 4 Jun, 1998.
Structural Symmetry and Model Checking – CAV 98, Vancouver, Canada – 29 Jun, 1998.
Enhancements to DICE – Intel Research, Portland, OR – Aug 1996.

PUBLICATIONS (DATA STREAMS)

B D VO AND G S MANKU
RadixZip: Linear Time Compression of Token Streams
VLDB 2007 (33rd International Conference on Very Large Data Bases), p. 1162-1172, Sep 2007.

G S MANKU, A JAIN AND A D SARMA
Detecting Near-Duplicates for Web Crawling
WWW 2007 (17th Annual World Wide Web Conference), p. 141-149, May 2007.

A ARASU AND G S MANKU
Approximate Counts and Quantiles over Sliding Windows
PODS 2004 (22nd ACM Symposium on Principles of Database Systems), p. 286-296, Jun 2004.

R MOTWANI, J WIDOM AND EIGHT PH.D. STUDENTS
Query Processing, Resource Management and Approximation in a Data Stream Management System
CIDR 2003 (1st Biennial Conference on Innovative Data Systems Research), p. , Mar 2003.

G S MANKU AND R MOTWANI
Approximate Frequency Counts over Data Streams
VLDB 2002 (28th International Conference on Very Large Data Bases), p. 346-357, Aug 2002.

G S MANKU, S RAJAGOPALAN AND B G LINDSAY
Random Sampling Techniques for Space Efficient Online Computation of Order Statistics of Large Datasets
SIGMOD 1999 (ACM International Conference on Management of Data), p. 251-262, Jun 1999.

G S MANKU, S RAJAGOPALAN AND B G LINDSAY
Approximate Medians and other Quantiles in One Pass and with Limited Memory
SIGMOD 1998 (ACM International Conference on Management of Data), p. 426-435, Jun 1998.

PUBLICATIONS (PEER TO PEER SYSTEMS)

I ABRAHAM, D MALKHI AND G S MANKU
(Brief Announcement) Papillon: Greedy Routing in Rings
DISC 2005 (19th International Symposium on Distributed Computing), p. 514-515, Sep 2005.

K KENTHAPADI AND G S MANKU
Decentralized Algorithms using both Local and Random Probes for P2P Load Balancing
SPAA 2005 (17th ACM Symposium on Parallelism in Algorithms and Architectures), p. 135-144, Jul 2005.

G S MANKU
Balanced Binary Trees for ID Management and Load Balance in Distributed Hash Tables
PODC 2004 (23rd ACM Symposium on Principles of Distributed Computing), p. 197-205, Jul 2004.

G S MANKU, M NAOR AND U WIEDER
Know thy Neighbor's Neighbor: The Power of Lookahead in Randomized P2P Networks
STOC 2004 (36th ACM Symposium On Theory of Computing), p. 54-63, Jun 2004.

P GANESAN AND G S MANKU
Optimal Routing in Chord
SODA 2004 (15th Annual ACM-SIAM Symposium on Discrete Algorithms), p. 169-178, Jan 2004.

G S MANKU
Routing Networks for Distributed Hash Tables
PODC 2003 (22nd ACM Symposium on Principles of Distributed Computing), p. 133-142, Jun 2003.

M BAWA, G S MANKU AND P RAGHAVAN
SETS: Search Enhanced by Topic Segmentation
SIGIR 2003 (26th Annual International ACM SIGIR Conference), p. 306-313, Jul 2003.

G S MANKU, M BAWA AND P RAGHAVAN
Symphony: Distributed Hashing in a Small World
USITS 2003 (4th Usenix Symposium on Internet Technologies and Systems), p. 127-140, Mar 2003.

PUBLICATIONS (MISCELLANEOUS)

G S MANKU AND J SAWADA
A Loopless Gray Code for Minimal Signed-Binary Representations
ESA 2005 (13th Annual European Symposium on Algorithms), p. 438-447, Oct 2005.

G S MANKU, R HOJATI AND R K BRAYTON
Structural Symmetries and Model Checking
CAV 1998 (10th International Conference on Computer Aided Verification, LCNS 1427), p. 159-171, Jul 1998.

S D GRIBBLE, G S MANKU, E A BREWER, T J GIBSON AND E L MILLER
Self-Similarity in File System Traffic
SIGMETRICS 1998 (ACM Conference on Measurement and Modeling of Computer Systems), p. 141-150, Jun 1998.

G S MANKU, M R PRASAD AND D A PATTERSON
Design and Analysis of a Voting Data Prefetch Engine
HiPC 1997 (4th International Conference on High Performance Computing), p. 100-105, Dec 1997.

G S MANKU, P JAIN, A AGGARWAL, L KUMAR AND S BANERJEE
Object Tracking using Affine Structure for Point Correspondences
CVPR 1997 (IEEE Conference for Computer Vision and Pattern Recognition), p. 704-709, Jun 1997.

PUBLICATIONS (MISCELLANEOUS)

G S MANKU
A Linear Time Algorithm for the Bottleneck Biconnected Spanning Subgraph Problem
Information Processing Letters (Vol 59, Number 1), p. 1-7, 8 July, 1996.

G S MANKU, A KUMAR AND S KUMAR
Circuit Partitioning with Partial Order for Mixed Simulation Emulation Environment
RSP 1995 (6th IEEE International Conferene on Rapid System Prototyping), p. 201-207, June 1995.

Last updated on September 19, 2007