

Ashish Gupta

Academic Background

Northwestern University, Evanston, IL

PhD Candidate in Department of Computer Science

July 2004- Early 2007

PhD Adviser: Prof. Peter Dinda

Master of Science in Computer Science

September 2002-June 2004

GPA: **4.00/4.00**

Indian Institute of Technology, New Delhi, India (IIT Delhi)

Bachelor of Technology in Computer Science and Engineering, July 1998 - May 2002.

Grade Point Average (CGPA): **9.04/10.00**

Thesis Title: QoS Aware Path Protection Schemes for MPLS Networks, Advisor: Prof. B.N. Jain

Thesis Awards:

Best Undergraduate Project Award in Computer Science 2001-2002

ICIM Stay Ahead Award for the best project in Information Technology and Computing Sciences 2001-2002

- XII Class Majors (CBSE Board) at DAV College, Chandigarh, India - Percentage 87.5%
- X Class Majors (ICSE Board) at Yadavindra Public School, Punjab, India - Percentage 89.2%

Refereed Publications

Research Interests: Distributed Systems, Networking and Security

- Robert Schweller, Yan Chen, Elliot Parsons, Ashish Gupta, Gokhan Memik and Yin Zhang, *Reverse Hashing for Change Detection on Highspeed Networks: Algorithms, Evaluation, and Applications*, Under submission to IEEE International Conference on Network Protocols 2005 (ICNP 2005)
- Ashish Gupta, Ananth Sundararaj, Marcia Zangrilli, Peter Dinda, Bruce B. Lowekamp, Free Network Measurement For Adaptive Virtualized Distributed Computing, Under submission to Supercomputing 2005
- A. Sundararaj, A. Gupta, P. Dinda, *Increasing Application Performance in Virtual Environments through Run-time Inference and Adaptation*, In Proceedings of the 14th IEEE International Symposium on High Performance Distributed Computing (HPDC 2005)
- Robert Schweller, Ashish Gupta, Elliot Parsons, Yan Chen, *Reverse Hashing for Sketch-based Change Detection*, In Proceedings of ACM SIGCOMM Internet Measurement Conference, October 2004, short paper version, Taormina, Sicily
- A. Sundararaj, A. Gupta, P. Dinda, *Dynamic Topology Adaptation of Virtual Networks of Virtual Machines*, Proceedings of the 7th Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers, Houston, 2004
- A. Gupta, P. Dinda, *Inferring the Topology and Traffic Load of Parallel Programs Running In a Virtual Machine Environment*, In Proceedings of the 10th Workshop on Job Scheduling Policies for Parallel Processing, June, 2004
- A. Gupta, B. Lin, P. Dinda, *Measuring And Understanding User Comfort With Resource Borrowing*, In Proceedings of the 13th IEEE International Symposium on High Performance Distributed Computing (HPDC 2004), Honolulu, Hawaii
- Ashish Gupta, Ashish Gupta, B.N. Jain, Satish Tripathi, *QoS Aware Path Protection Schemes for MPLS Networks*, In Proceedings of International Conference on Computer Communications, August 2002 (ICCC 2002)

Accepted Posters (New projects in progress)

- q Ashish Gupta, Manan Sanghi, Peter Dinda, Fabian Bustamante, *Magnolia: a novel DHT architecture for Keyword based search*, Poster Presentation at Network System Design and Implementation (NSDI 2005), Boston
- q Ashish Gupta, Peter Dinda, Fabian Bustamante, *Distributed Popularity Indices*, Poster Presentation at ACM SIGCOMM 2005, Philadelphia

Honors and Awards

Programming Awards

- § Represented **India and Computer Science Department, IIT Delhi** at **ACM-IBM International Collegiate Programming Contest** World Finals held at Vancouver, Canada, and achieved **29th position** among 2,200 competing teams from 70 countries, March 2001
URL : <http://icpc.baylor.edu/past/icpc2001/Finals/Standings.html>
- § **2nd position** among 70 teams in South Asia Regionals of **ACM-IBM International Collegiate Programming Contest** held at IIT Kanpur and were selected for the World Final, Dec 2000
URL: <http://www.cse.iitk.ac.in/users/acm/archive/acm2000/acm2000/Result.html>
- § **Programming Excellence Award** by *Upsilon Pi Epsilon* International Honor Society for Computing Sciences for participation in the ACM-IBM ICPC 2001 World Finals, March 2001
- § **2nd position** in Bitwise – an International Algorithm Intensive Programming Contest organized by IIT Kharagpur, amongst 300 teams, Feb 2002
URL : <http://www.iitkgp.ernet.in/~cse/bitwise>
- § **1st prize twice** (during 1st year and 3rd year) in 'Software Programming Contest' in the All India Technical Festival (TRYST) organized by IIT Delhi (Feb 1999 and Feb 2002)
- § Reached National levels twice (**amongst top 8 in India**) of All India High School Computer Programming Contest organized by **Computer Society of India** in class IX and X. (Feb 1995 and Feb 1996)

Research Related

- § Awarded **Best Undergraduate Project Award** in the Department of Computer Science, IIT Delhi for our Final Year Project "QoS Aware Path Protection Schemes in MPLS Networks", May 2002
- § Awarded the **ICIM Stay Ahead Award** for the best Undergraduate Final Year Project, IIT Delhi in Information Technology and Computing Sciences among Undergraduates, March 2002
- § Selected for the **Summer Undergraduate Research Award** after 2nd year at IIT. This is given to selected project proposals from various departments for encouraging research at UG level after a highly selective process, May 2000
- § Nominated for the IBM PhD fellowship by the Computer Science Department, Northwestern University for 2005-06;
- § Selected for the **Student Volunteer Program** at **Supercomputing 2004**, the premier supercomputing expo and conference, to be held in November 2004, Pittsburgh. Involved in technical infrastructure setup.

Academic Awards

- § **All India Rank 10th** among 130,000 students in JEE '98 (Joint Entrance Examination) for admission into **Indian Institutes of Technology (IITs)**
- § **All India Rank 17th** in Roorkee Engineering Exam, **3rd Rank** in Common Engineering Test. and **9th Rank** in Punjab Engineering Test (National Level Engineering Exams in India, conducted after High School)
- § Awarded the **IIT Merit Award 3 times** for being among **top 7%** in **Department of Computer Science, IIT Delhi** in 3rd year (**2nd position**) and 4th year (**1st position**).
- § Received **1st prize** for topping D.A.V. College in High School, XI class (batch of 400 students)
- § Air Force Trophy for the "**Best All Round Student**" in School for excellence in academics and extra-curricular activities
- § **Co-curricular**: Numerous prizes during my schooling and college, such as English writing, Quizzing, Special Distinction in Math and Computers.

Research and Professional Experience

Current Research Areas

My current research looks at intelligent and automated resource adaptation for improved performance and resource allocation in large scale distributed applications and Grid services, especially for interactive purposes; We utilize automated network and application topology inference and traffic measurement; currently actively involved with Virtual Machine abstractions and virtual overlay networks for wide area distributed computing.

I am also involved in finding ways of making intelligent and more aggressive use of desktop resources without irritating the user, for background apps like Google Compute, Seti@Home and Folding@Home (see our User Comfort Project).

Besides, I am also involved in the Network Security area and data stream computation, especially in finding and reporting anomalies in very high-speed traffic streams (10-20 Gbps links) in real-time.

On the development side, I am interested in developing innovative web applications, which elevate the user's web and shopping experience e.g. intelligent shopping assistants (through web-wide aggregation)

Key words: Distributed computing, Grid computing, Virtual machines, virtual networking, operating systems, high performance computing and networking, performance analysis and prediction, network measurement, network and application inference, resource borrowing, user behavior, network security, traffic anomalies, web applications, web services, gossip communication, real time systems

IBM Research Labs, New Delhi, India (Summer 2001)

With **Dr. P.V. Kamesam**, Research Scientist, IBM Watson Research Lab

Content Management and Distributed Application Development in context of eGovernance & Security/Audit Management in Distributed Databases

Using IBM's Content Manager, I developed a new C++ API and middleware for special distributed data storage needs of eGov applications (distributed storage and replication) and then developed a prototype eGov application on top of it for national voter registration.

I then worked on the problem of managing Security/Audit Information in Distributed Database Systems (like eGovernment) and our team suggested a new scalable Architecture for the above problem and also worked on other issues like secure transfer and storage.

Research Assistantship (2002-now)

Research Assistantship with Prof. Peter Dinda, Computer Science Department, Northwestern University

Virtuoso

URL: <http://www.cs.northwestern.edu/~plab/Virtuoso>

Imagine ordering a bunch of machines (or even a cluster) from a website and just start using them with a single click, within your browser! Virtuoso aims to leverage Virtual Machines and the distributed computing mechanisms to provide a new abstraction for computing. This infrastructure offers us unique opportunities for performance adaptation according to user's needs. I have worked on automated application inference, which can automatically figure out the network/computational needs of a distributed applications and then automatically adapt the underlying network and virtual machine placement to optimize performance and cost for the user as well as the providers. We have demonstrated that we can significantly boost application performance without any user intervention.

Understanding and Measuring User Comfort with Resource Borrowing

URL: <http://comfort.cs.northwestern.edu>

How much of CPU, Memory and disk can be borrow (for useful computation like Folding@Home or Google Compute) from a typical machine without irritating the user ? Our results suggest that more aggressive resource usage is possible, not just in screen saver mode, but all the time. This has many other interesting applications.. The results of our work could be applied in many different ways:

- Ø to enable applications like SETI@HOME and Folding@HOME to borrow resources more aggressively without invading users' comfort zone,
- Ø understanding resource requirements for several virtual machines located on a single host so that they feel like full machines,
- Ø inspire design of new interactivity-aware scheduling algorithms,
- Ø understand how to provide cheaper computing while ensuring user satisfaction

Teaching Assistant – Real Time Systems and Sensor Networks (Winter 2005)

As a teaching assistant and developer, I am actively involved in development of a sensor network infrastructure with Pocket PC and Berkley Motes, which talk to each other to convey real time sensor, audio and video information to desired locations and the web. I am involved in designing and developing protocols, interfaces and necessary middleware to support such applications. I've had significant experience in programming over Bluetooth, Zigbee, WiFi, audio/video programming and streaming, and heterogeneous platform message transfers. Our goal is to design a complete building monitoring system using sensors and Pocket PCs.

Online Detecting Network Anomalies on very High speed links using Sketch Techniques

With Prof. Yan Chen, Northwestern University

This project looks at finding anomalies in very high-speed traffic streams (10-20 Gbps links) in real-time. We propose algorithms to detect heavy change in network traffic without the need of any off line analysis of traffic. These algorithms are based on a useful data structure, the k-ary sketch

B. Tech. Thesis: QoS Aware Path Protection Schemes for MPLS Networks (2001-2002)

http://www.cs.northwestern.edu/~agupta/_projects/mpls_pp/

With Prof. B.N. Jain, Department of Computer Science & Engineering, IIT Delhi

Multiprotocol Label Switching integrates Layer 2 information about network links (bandwidth, latency, and utilization) into Layer 3 (IP) in order to simplify and improve IP-packet exchange. The problem of path protection looks at fast rerouting of critical traffic along pre-allocated paths in case of a node/link failure. Our focus has been to develop algorithms to provide QoS guarantees even in case of failure along with efficient bandwidth reservation. We considered a new approach based on dividing the entire route into segments and protecting each segment separately. We have developed algorithms for providing bounded switch over time after failure, end-to-end delay, jitter, their combinations and bound on reliability. We have tested the algorithms in a self developed LSP Admission Control Simulator (in C++) and the algorithms perform very well compared to existing protection schemes, utilizing lesser bandwidth.

This project received the "Best B.Tech Project Award" in CS Department and also the ICIM Stay Ahead Award for the best project in Information Technology and Computing Sciences.

The results were also published in the International Conference of Computer Communications, 2002.

Summer Undergraduate Research Award for project titled "Optimal Resource Utilization in VSAT networks" (Summers 2000)

http://www.cs.northwestern.edu/~agupta/_projects/visual_mac_simulator

With **Prof. H.M. Gupta**, Electrical Dept. IIT Delhi

In this project our aim was to provide optimal resource utilization for various traffic streams in VSAT networks. We looked at various types of traffic such as bursty, demand based traffic and multimedia traffic and worked on the problem of optimal utilization of the satellite bandwidth under various traffic conditions. We also developed a GUI oriented discrete-event simulator in Visual C++ and implemented many MAC layer protocols like ALOHA, R-ALOHA, FDMA, TDMA, frequency hopping and their combinations for simulation and study of various scenarios.

System Administrator of Computer Lab at IIT Delhi (2000-2001)

I was selected for the position of **Computer Secretary** (a voluntary position) of my hostel and was the **System Administrator** of the 24 Hr Computer Lab for one year in my student dorm, IIT Delhi in 3rd year, which had over 400 users. The position demands lot of interaction with the students and maintaining orderly conduct of machines as well as the users round the clock.

Teaching Assistantships

- q Computer Networks (involved in teaching)
- q Real Time Systems and Sensor Networks (involved in teaching)

Technical Skills

- § Proficiency in C, C++, Java (Swing, Network Programming), Assembly and familiarity with VHDL, Scheme, SML, Pascal, Perl/CGI and shell programming
- § Well Experienced in Visual C++/Microsoft Foundation Classes (MFC), Visual Basic, Borland C++ Builder, Borland JBuilder, Borland Delphi
- § Proficiency in Simulator Design, Systems Programming, Socket Programming, GUI Development (in many environments), Database Programming, 3D Development/OpenGL Programming, Image Processing
- § Experience in XILINX hardware development tools and VHDL.
- § Parallel computing: PVM Programming
- § Experience in NS, Latex, GNU Plot.
- § Experience and interest in System and Graphics Programming including Game programming

Mobile and Sensor Platforms

Pocket PC 2003 (Audio, Video, Bluetooth and Wireless Programming), Experience with MicaZ sensor nodes programming (TinyOS, nesC programming), and intercommunication between micaZ and Pocket PC devices

Mathematical skills:

Statistical characterization, analysis and prediction, Probability and Stochastic Models, Queuing Theory, Linear and Non-Linear Optimization

Communication Skills

- § Was the Teaching Assistant for Computer Networking course and was involving in teaching the class. Rated as Excellent in the Course Evaluation.
- § Member of Toastmasters Club, Evanston, IL
- § Did the *Giving effective research talks* course, which focuses on giving short talks that are interesting and motivating. Conducted by **Prof. Don Norman**, Northwestern University.
- § Did the *Technical Communications* Course, which focuses on English Technical writing skills, presentation techniques and technical report writing. I was awarded an "A" grade in this course. My oral presentation was rated as one of the best in the class (18/20).
- § My oral presentation on "*DNA Computing*" in the Colloquium course received the highest rating. (45/50). For the slides, see **Work Samples** above

Past Major Research/Development Projects

My research and development experience briefly: (Most publications and software is available for download from <http://www.cs.northwestern.edu/~agupta>)

Computer Science Area	Projects and Software
Systems/Networking	Virtuoso Understanding User Comfort w/ Resource Borrowing, Online Anomaly Detection in very high speed networks Fast Resource Dissemination using Gossip Models in P2P Networks QoS Aware Path Protection in QoS Networks Optimal Resource Util. in VSAT Networks
Architecture/Hardware	Analysis and Visualization for Automated Synthesis of Embedded Systems Hardware Level SPARC Simulator SPARC Assembler Design/Implementation of 4 Port Packet Switch
Graphics	Lights! Camera! Action! : 3D Rendering and Animation Fractal Image Compression Space Invaders 3D: OpenGL implementation Image Morphing using Delaunay Triangulation
AI	An Intelligent Agent for Chinese Checkers using minmax
Algorithms/Theory	Optima: Tool for Linear Optimization Study of Evolutionary Algorithms Visual Graph Studio with Graph Algorithms
Distributed Databases	Distributed Content Management and Application for eGovernance

Project Descriptions

Fast Resource Information Dissemination in a P2P Network using Gossip Models (Winter 2003)

With Prof. Fabian Bustamante, Northwestern University.

With the goal of designing an efficient P2P resource information dissemination system, we have focused on finding an efficient communication model for propagation of this information with **network resource scalability** and **freshness of information** as two primary concerns. We studied various combinations of dynamic gossip schemes using simulation (including a rank based spatial gossip scheme called bin-halving) and found that depending on the metric, different schemes may be preferred. One of the main parts of our work was design of meaningful metrics for representing "freshness" and then evaluating various schemes with respect to these.

Design, Simulation and Implementation of a simple 4 port Packet Switch (Spring 2000)

With Prof. Balakrishnan, CSE Department, IIT Delhi

This project involved design and implementation of a prototype packet switch. The switch had the ability to switch packets between various ports based on packet **destination address**. We designed the Switch in Xilinx Design Tools and simulated it and also wrote the Behavioral and Structural Description of the Switch in VHDL. Then we finally implemented the Switch on FPGA.

Discrete-Event Hardware Level SPARC Simulator with a two pass Assembler (Fall 1999)

http://www.cs.northwestern.edu/~agupta/_projects/sparc_simulator

With Prof. Anshul Kumar, CSE Dept. IIT Delhi

This project involved understanding the working of the SPARC processor and then developing a simulator for its data path. We developed an event-driven hardware block level simulator for the SPARC chip, which also supports instruction pipelining. It also allows timing simulation taking into account the various delays in the datapath and aids in determining the optimal clock speed. A two pass Assembler was also written to convert SPARC assembly to SPARC machine code. The simulator has a comprehensive GUI written in Visual C++ to illustrate what happens inside the processor at each step of execution (for educational purposes)

Integrated Framework for Analysis and Visualization for Automated Synthesis of Embedded Systems (Fall 2000)

URL: <http://www.cse.iitd.ernet.in/esproject>

With Prof. Balakrishnan and Prof. Anshul Kumar, CSE Dept. IIT Delhi (Mini Project, 2000)

The project involved development of a Java-based tool to assist in the development of embedded systems. The user should be able to analyze and visualize his design of the embedded system and hence come up with the optimum partitioning of the tasks between hardware and software. The tool also allowed automatic hardware-software partitioning. It allows specification of the various components of the embedded system like processors, memory, cache, bus, ASICs and so on and their interconnections. Further, it allows analysis of the entire system to test the performance of the target platform configuration. It also provides visualization of the performance of the resultant system as well as the overall architecture.

Lights, Camera, Action! A 3D Rendering and Animation Package in Java (Spring 2001)

http://www.cs.northwestern.edu/~agupta/_projects/3d_rendering/Web/index.html

With Prof. K.K. Biswas, CSE Dept., IIT Delhi

This project involved design of general purpose 3D software which illustrates various concepts learnt in Computer Graphics course. It allows rendering and animation of a scene consisting of multiple 3D objects, features a complete scripting language for scene description and animation control, includes lots of primitive objects like cylinder, sphere, cone, ellipsoid, donut, super-ellipsoids etc. with wire-frame based modeling. We also implemented various lighting models like Diffusion and Specular models. We also implemented 3D fractals like trees and rough terrain modeling. Animation includes straight-line motion as well curvilinear motion and supports animation of multiple objects. It also allows animation of light sources to provide interesting lighting effects to the scene.

URL: http://www.cs.northwestern.edu/~agupta/_projects/3d_rendering/Web/

Study and Implementation of Fractal Image Compression (Spring 2001)

http://www.cs.northwestern.edu/~agupta/_projects/image_processing/web/FractalImageCompression/

Under guidance of Prof. Prem Kalra, CSE Dept, IIT Delhi

In this project we studied the technique of fractal image compression and implemented it in C++. We carried out various analyses to study its efficiency and advantages and also compared it to JPEG.

Artificial Intelligence Project: An Intelligent Chinese Checkers playing program (Fall 2001)

http://www.cs.northwestern.edu/~agupta/_projects/chinese_checkers/web

With Prof. K.K. Biswas, CSE Dept., IIT Delhi

In this project we developed an Intelligent Chinese Checkers playing program. We developed five heuristics, which made it almost impossible to beat the computer player. We also studied and did some optimizations in the Minimax algorithm which greatly improved the speed of the computer player. We were able to greatly improve alpha-beta pruning in the Minimax algorithm.

Other projects

- § Discrete-Event Simulator in Java to simulate queues (M/M/1 and other) for analysis based on Queuing Theory.
- § Client/Server communication software in Linux involving Berkley sockets and I/O Multiplexing.
- § Multi-threaded Simulator for "1 bit Sliding Window" Protocol with analysis in Java.
- § "Optima" - a tool for linear optimization using the Simplex Algorithm in Visual C++.
- § Visual Graph Editor, which allows for point-n-click construction of directed graphs with export option to files (in Data Structures course).
- § Complete database program with implementation of indexes and relations in C++/OOP.
- § Studied and implemented the JPEG Digital Image Compression Format in C++.
- § Image Morphing based on 2-D triangulation method in C++ as a part of Digital Image Processing.
- § Implementation and Study of Fibonacci and Binomial Heaps as part of Algorithm Design/Analysis.

Mathematical and Software Skills

Administrative Experience

- § Was elected as the Class Convener during 2nd year and was responsible for facilitating teacher-student

interactions.

- § Organized a trekking trip to McLeod Ganj (the home of Buddhists in India) in Himachal Pradesh for the CSE Department in Feb 2000.

Teaching Experience

Teaching Assistant for Computer Networks course, Winters 2004

Have taught many students Physics and Mathematics for IIT JEE engineering entrance examination.

Have taught many company employees and students in developing computer and programming skills

Memberships

Student Member of ACM and ACM SIGCOMM

References

Prof. Peter Dinda

Assistant Professor

Computer Science Department

Northwestern University

Phone: (847) 467-7859

Email: pdinda@cs.northwestern.edu

Prof. Ming-Yang Kao

Department Chair

Computer Science Department

Northwestern University

Phone: (847) 563-0426

Email: kao@cs.northwestern.edu

Prof. Fabián E. Bustamante

Assistant Professor

Computer Science Department

Northwestern University

Phone: (847) 491-2745

Email: fabiánb@cs.northwestern.edu

Prof. Yan Chen

Assistant Professor

Computer Science Department

Northwestern University

Phone: (847) 491-4946

Email: ychen@cs.northwestern.edu

Courses Done**Graduate**

Non-Linear Optimization
Real Time Systems
Advanced Statistics
Network Security: A Measurement based approach
Analysis and Prediction of the Dynamic Behavior of Applications, Hosts, and Networks
Advanced Operating Systems
Resource Virtualization
Distributed Computing
Introduction to Networking
Operating Systems
Computer Graphics
Giving Effective Research Talks

Undergraduate**Theory and Mathematics**

Introduction to Computer Science
Data Structures and Algorithms
Discrete Structures
Programming Languages
Introduction to Logic for Computer Science
Theory of Computation
Algorithm Design and Analysis
Optimization - Methods and Applications
Introduction to Analysis and Differential Equations
Mathematics II - Complex Analysis, Linear Algebra, Fourier Series
Probability & Stochastic Processes

Systems and Networks

File Structures and Systems
Database Management Systems
Operating Systems
Computer Architecture
Advanced Computer Networks
Computer Networks – I

Computer Engineering

Introduction to Electronic Circuits
Digital Electronics
Electronics Laboratory I
Signals and Systems
Computer Architecture
Digital Hardware Design
Digital Hardware Design Lab
Mini Project (Embedded Systems)

Other

Artificial Intelligence
Computer Graphics
Digital Image Processing
Technical Communications