

Gaurav Khanna
502 Apt 10 Stinchcomb Dr.
Columbus, OH 43202, USA

614-804-8765
khannag@cse.ohio-state.edu
Web: <http://www.cse.ohio-state.edu/~khannag>

Objective

To obtain a research oriented challenging full-time position in the field of High Performance Data-Intensive Computing with special emphasis on Job Scheduling and Resource Management in Clusters and Grids, Performance Modeling and Parallel I/O systems.

Education

- **The Ohio State University** Columbus, OH
Ph.D., Computer Science and Engineering (GPA: 3.7/4.0) Sep. 2003 - Sep. 2008
 - **Dissertation Title:** A Data-Locality Aware Mapping and Scheduling Framework for Data-Intensive Computing
- **Delhi Institute of Technology, University of Delhi** Delhi, India
Bachelor of Engineering, Computer Science and Engineering 1998-2002
 - Ranked 3rd in a class of 65.

Research Interests

Scheduling algorithms, Data-Intensive Computing, Parallel I/O.

Work Experience

- **The Ohio State University** Columbus, OH
Graduate Research Associate Apr 2004 till date
 - **Data Locality Aware Job scheduler for Clusters**
Designed and implemented locality conscious job mapping algorithms for data-intensive computing on homogeneous clusters. The objective is to exploit data reuse to yield load balanced I/O volume minimizing mapping of jobs onto compute nodes. I have extended this work for job mapping on systems consisting of multiple heterogeneous clusters. These works have been published at CCgrid05 and HCW06.
 - **Coordinated Job Scheduling and Data Replication**
Designed and implemented algorithms for coordinated scheduling and data replication using both coupled and decoupled approaches. In this work, I have studied the interaction between job scheduling and data replication and designed various approaches to solve this problem. This work was published at HPDC06.
 - **Efficient Data Mover for Data-Intensive computing**
Designed and implemented algorithms for scheduling of file transfers on heterogeneous systems. These works have been accepted at Europar07 and IPDPS08.
 - **Analysis of use of PVFS for Jobs exhibiting Pipelined Shared I/O**
Investigated the impact of striping unit size and the number of stripes on the performance of pipeline shared I/O jobs over PVFS. This work was accepted at Grid04.

- MCS Division, Argonne National Laboratory** IL

 - *Summer Intern* *June 2007 - Sep. 2007*

 - **Split-TCP for GridFTP**

Developed an enhanced DSI interface for GridFTP which enables splitting of end-to-end TCP connections resulting in improved performance. Also, developed a GridFTP client using Globus C-WS core which can allocate dedicated bandwidth for a wide-area transfer over high-speed network links.
- IBM, India Research Lab** Delhi, India

 - *Summer Intern* *June 2006 - Sep. 2006*

 - **Parallel Job Mapping on BlueGene/L**

Designed and implemented heuristics for addressing the parallel job mapping problem on BlueGene/L. The proposed heuristics combined graph partitioning and embedding techniques to generate efficient job mappings.
- C-DOT (Center for Development of Telematics)** Delhi, India

 - *Research Engineer* *Sep. 2002 - July 2003*

 - Designed and implemented GTP (GPRS tunneling protocol) in C language and integrated it with the interacting protocol layers
 - Worked in a team to develop user domain protocol stacks for the 3G-UMTS network
- AmSoft Systems India Inc.** Delhi, India

 - *Summer Intern* *May 2001 - Aug 2001*

 - Developed a library in VC++ to generate 3D bar charts, 3D pie charts, scatter charts and line charts for input data which could be either integer or floating point. The DLL was integrated into one of AmSoft's ongoing projects for the purpose of data analysis.
- The Ohio State University** Columbus, OH

 - *Graduate Teaching Associate* *Sep. 2003 - March 2004*

 - Taught an Introductory level course in Computers (CSE100) to a class of 40 students. The course is an intensive introduction to problem solving using business application software.

Publications

- Multi-Hop Path Splitting and Multi-Pathing Optimizations for Data Transfers over Shared Wide-Area Networks using GridFTP (Short Paper).
Gaurav Khanna, Umit Catalyurek, Tahsin Kurc, Raj Kettimuthu, P. Sadayappan, I. Foster, Joel Saltz.
17th IEEE International Symposium on High-Performance Distributed Computing, (HPDC 2008)
- A Dynamic Scheduling Approach for Coordinated Wide-Area Data Transfers using GridFTP.
Gaurav Khanna, Umit Catalyurek, Tahsin Kurc, Raj Kettimuthu, P. Sadayappan, Joel Saltz.
22th IEEE International Parallel and Distributed Processing Symposium, (IPDPS 2008),
Acceptance rate: 25.5%
- Scheduling File Transfers for Data-Intensive Jobs on Heterogeneous Clusters.
Gaurav Khanna, Umit Catalyurek, Tahsin Kurc, P. Sadayappan, Joel Saltz.
13th European Conference on Parallel and Distributed Computing, (Europar 2007), *Acceptance rate: 26.5%*
- Globus GridFTP: What's New in 2007 (Invited Paper).
John Bresnahan, Mike Link, **Gaurav Khanna**, Z. Imani, Raj Kettimuthu, Ian Foster.
1st International Conference on Networks for Grid Applications, (GridNets 2007)

- Task Scheduling and File Replication for Data-Intensive Jobs with Batch-Shared I/O.
Gaurav Khanna, Nagavijayalakshmi Vydyanathan, Umit Catalyurek, Tahsin Kurc, Sriram Krishnamoorthy, P. Sadayappan, Joel Saltz.
*15th IEEE International Symposium on High Performance Distributed Computing, (HPDC 2006),
Acceptance rate: 15%*
- A Data Locality Aware Online Scheduling Approach for I/O-Intensive Jobs with File sharing.
Gaurav Khanna, Umit Catalyurek, Tahsin Kurc, P. Sadayappan, Joel Saltz.
12th Workshop on Job Scheduling Strategies for Parallel Processing, (JSSPP 2006)
- A Hypergraph Partitioning Based Approach for Scheduling of Tasks with Batch-shared I/O.
Gaurav Khanna, Nagavijayalakshmi Vydyanathan, Tahsin Kurc, Umit Catalyurek, Pete Wyckoff, Joel Saltz, P. Sadayappan.
5th IEEE/ACM International Symposium on Cluster Computing and the Grid, (CCGrid 2005)
- Use of PVFS for Efficient Execution of Jobs with Pipeline-Shared I/O.
Nagavijayalakshmi Vydyanathan, **Gaurav Khanna**, Tahsin Kurc, Umit Catalyurek, Pete Wyckoff, Joel Saltz, P. Sadayappan.
5th IEEE/ACM International Workshop on Grid Computing (Grid 2004) - In Conjunction with SuperComputing 2004
- Scheduling of tasks with Batch-Shared I/O on Heterogeneous Systems.
Nagavijayalakshmi Vydyanathan, **Gaurav Khanna**, Umit Catalyurek, Tahsin Kurc, P. Sadayappan, Joel Saltz.
15th Heterogeneous Computing Workshop (HCW 2006) - In Conjunction with IPDPS 06.

Presentations

- Task Scheduling and File Replication for Data-Intensive Jobs with Batch-Shared I/O. At the 15th IEEE International Symposium on High Performance Distributed Computing, (HPDC 2006), Paris, France, June 2006.
- Scheduling File Transfers for Data-Intensive Jobs on Heterogeneous Clusters . At the 13th International Euro-Par Conference European Conference on Parallel and Distributed Computing, (Europar 2007), Rennes, France, August 2007.

School Projects

- Designed and implemented a simulator for analyzing the performance of a data-intensive parallel application on a cluster of workstations.
- Implemented an interpreter for Prolog in the ML language.
- Time Series Data Miner: Designed and implemented a frequent itemset miner for time series data.
- Implemented a simulator for dynamic instruction scheduling using Tomasulo's algorithm
- Implemented a simulator for network congestion control using the Leaky Bucket algorithm.

Technical Skills

Languages: C, C++, Java, VB, VC++, Perl, ML, Parallel Programming (MPI, OpenMP).

Operating Systems: Linux, Solaris, UNIX, Windows 95/98/NT/2000/XP.

Applications: Matlab, Debuggers (gdb), Configuration Management (CVS, RCS), Microsoft Office.

Academic Background: Advanced Computer Architecture, Advanced Operating Systems, High Performance Computing, Algorithms, Parallel Computing, Distributed Systems, Data Mining, Compilers/Programming Languages and Databases, Computer Networking

Awards

Graduate Teaching Assistantship (GTA) at OSU	2003-2004
Graduate Research Assistantship (GRA) at OSU	2004-2007
Recipient of Institute Merit Scholarship and Tuition Fee Waiver, University of Delhi	1998-2002
Secured All-India Rank 1618 (top 1.1%) among 150,000 candidates in IIT-JEE	1998
Secured All-India Rank 103 (top 0.5%) among 50,000 candidates in Common Entrance Exam.	1998
Awarded Junior Science Talent Search Scholarship	1995
Travel Scholarship to attend HPDC conference	2006

Professional Activities

- Reviewer for International Conference on Distributed Computing Systems (ICDCS) 2005.
- Reviewer for International Parallel and Distributed Processing Symposium (IPDPS) 2007.

References

- Available upon request.