



# Dong Lu

Department of Computer Science  
Northwestern University  
1890 Maple Ave. Suite 300  
Evanston, IL 60201, USA

(847) 467-1006 (Office)  
(312) 363-9894 (Cell)  
donglu@cs.northwestern.edu  
<http://www.cs.northwestern.edu/~donglu>

## *Profile*

Highly motivated and experienced system researcher, designer and developer; Resourcefulness in developing innovative solutions; Well educated for interdisciplinary tasks; Strong team spirit and leadership.

## *Research Interests: Distributed Systems and Networking*

Grid computing; Empirical queuing models for computer systems; Peer-to-Peer systems; End system multicast; Performance analysis; Internet measurement, statistical characterization and prediction; High performance networking; Wireless Ad-Hoc networks.

## *Educational Background*

- **Northwestern University**, Evanston, IL, USA. (Sept. 2000 ~ Present)  
**PhD** in *Computer Science*, expected in May 2005 (**GPA 3.82**).  
**M.S.** in *Computer Science*, June 2002 (**GPA 3.92**).
- **Shanghai Jiao Tong (Chiao Tung) University**, Shanghai, China. (Sept. 1997 ~ Jan. 2000)  
**M.S.** in *Mechanical Engineering (National Die & Mold CAD Engineering Research Center)*, January 2000 (**GPA 3.70**).
- **Central South University**, Changsha, China. (Sept. 1993 ~ Jun. 1997)  
**B.S.** in *Material Science and Engineering*, June 1997 (**GPA 3.60**).

## *Employment History*

- **Microsoft Research Asia**, Wireless and Networking Group. (Jun. 2004 ~ Sept. 2004)  
*Summer Intern.* Research on the incentives in wireless Ad-Hoc networks. Designed a power-aware reputation system to support power-aware routing in non-cooperative ad-hoc networks.
- **Argonne National Laboratory**, Mathematics & Computer Science Division. (Jun. 2003 ~ Sept. 2003)  
*Summer Intern.* Research and development on the Quality of Service for Grid services. Designed and developed a QoS enabled Grid Service.
- **Northwestern University**, Computer Science Department. (Sept. 2000 ~ Present)  
*Teaching Assistant & Research Assistant.*
- **Alcatel China**, R&D Center. (Jan. 2000 ~ Jun. 2000)  
*R&D Software Engineer.* R&D on the Operation and Maintenance Center of the Alcatel GSM cell phone network.

## *Skills*

### ❖ **Programming Languages:**

- **Proficient with:** C, C++/STL, Perl (8 years of C++, above 82% in a recent **BrainBench C++** test).
- **Familiar with:** Java, C#, SQL, FORTRAN, Matlab, JavaScript, DHTML, XML.

### ❖ **Development Environments:**

- **Linux & UNIX:** POSIX multithread programming; Socket programming; PVM and MPI on parallel computers and Linux cluster; Globus Toolkit 3.0; ns2 network simulator; Linux Kernel Hacking on TCP congestion control; CVS; SOAP
- **Windows:** Visual C++; MFC; Winsock programming; Multithreads programming; GUI programming with Visual C++; OpenGL programming with Visual C++; Windows API programming with Borland C++ and C; Fundamentals of .NET and C#
- **Database:** Oracle9i C++ Call Interface; PL/SQL; Perl DBI programming with Oracle9i and MySQL; Some experience with Oracle9i DBA and Linux system administration; ODBC programming with Visual C++;

### ❖ **Mathematics**

- Statistics; Linear algebra; Empirical Queuing; Stochastic modeling; Time series models; Numerical analysis; Partial differential equations; Fundamentals of game theory;

## *Recent Projects*

- **Virtualized Audio System.** The idea behind virtualized audio is to extract sound sources from their native acoustic spaces and insert them into any virtual spaces. The system consists of a windows client and a cluster of Linux servers. The client communicates with servers at socket level, while the servers communicate with each other using PVM. (C++ on both Windows and Linux; GUI with MFC; multi-threading; socket level programming; OpenGL; parallel computing with PVM on a Linux cluster)
- **Power-aware reputation system for wireless Ad-Hoc networks.** We introduce power-aware routing into the non-cooperative Ad-Hoc networks with the support of a power-aware reputation system. We design and implemented the reputation system into the current network simulator 2 which consists of about 100,000 lines of C++ code. (VC++ on Windows XP)
- **Fat-tree based end-system multicast.** FatNemo is a novel scalable peer-to-peer multi-source multicast protocol based on the idea of fat-trees. For many-to-many multicast applications, such as video conferencing, this eliminates the bottlenecks inside the overlay network. Our simulations and wide area experiments show that FatNemo not only minimizes the average and the standard deviation of the response time, but also handles end host failures gracefully without suffering a performance penalty. (Java, multithreading, sockets level programming on Linux)
- **Looking at the server side of P2P systems.** We start by characterizing the peer server-side workload based on extensive trace collection and statistical analysis. We then designed scheduling schemes for the server side of the P2P file sharing systems. We built our schedulers into an open source P2P applications and evaluated its performance on the Internet. Our results show that average response time can be dramatically reduced by more effectively scheduling the requests on the server-side of P2P systems. (C++, multithreading, sockets level programming on Linux, Statistics)
- **Characterizing and predicting TCP throughput on the Wide Area Network.** We first built a TCP throughput measurement framework to conduct large-scale Internet measurement study. We first statistically characterized the transient End-to-End TCP throughput distribution, explored the strong correlation between TCP flow size and throughput. Then based on the observations, we proposed a novel yet simple TCP throughput prediction model with time-series predictors. We evaluated the prediction system on the Internet. (C, socket level programming; Statistics and Time series; Perl)
- **Parallel TCP throughput prediction on the Wide Area Network.** Based on ns2 simulations and experiments on the wide area network, we first studied how does parallel TCP gain performance and then set up a half-analytical, half-empirical prediction model to predict the throughput of parallel TCP. (C, socket level programming on Linux; Statistics; Perl)
- **Domain-based Scheduling on Web Servers.** We first hacked the logging module of Apache web server to collect accurate performance traces. Then we built a simulator to conduct empirical queuing studies. A new scheduling policy is proposed based on the observation and its effectiveness is verified with extensive simulation study. (C++, event driven simulations, queuing theory)
- **Quality of Service for Grid Services.** We implemented a prototype Grid service with prediction based, soft QoS guarantees. (Java, C, IPC on Linux; Time series)
- **The Unified Relational Grid Information Services Project.** This project seeks to apply the relational database model (Oracle 9i) to static and dynamic GRID information. Soft real-time queries were studied. Efforts are being taken to make the systems highly scalable and fault tolerant. (Perl, C++, Oracle 9i C++ Call Interface, PL/SQL)
- **GridG.** A Toolkit developed to generate realistic synthetic computational Grids, which are modeled as a reasonably annotated graph that follows the discovered power laws of Internet topology. Contributions to the networking theory were made by pointing out a contradiction between two power laws of Internet topology. The software was released online.
- **Linux kernel hacking on aggressive TCP.** The project is aimed at improving the performance of TCP/IP based cluster-computing applications without modifying any code of the application such as IBM DB2 and PVFS. (C, Linux kernel hacking)

## *Software Released*

- **GridG:** An extensible toolkit for generating synthetic Computational Grids. Released online at: <http://www.cs.northwestern.edu/~urgis/GridG>
- **TameParallelTCP:** Estimate throughput and impact of parallel TCP flows. Released online at: <http://www.cs.northwestern.edu/~plab/Clairvoyance/Tame.html>

## *Selected Refereed Publications* (Full publication list available in the CV)

- **Dong Lu, Yi Qiao, Peter Dinda, Fabian Bustamante,** "Characterizing and Predicting TCP Throughput on the Wide Area Network", In Proceedings of the 25th IEEE International Conference on Distributed Computing Systems (**ICDCS 2005**), June 2005, Columbus, Ohio, to appear.

- **Dong Lu**, Yi Qiao, Peter Dinda, Fabian Bustamante, "Modeling and Taming Parallel TCP on the Wide Area Network", In Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (**IPDPS 2005**), April, 2005, Denver, Colorado. To appear.
- **Dong Lu**, Haitao Wu, Qian Zhang, Wenwu Zhu, "PARS: Stimulating Cooperation for Power-Aware Routing in Ad-Hoc Networks", In Proceedings of the 40th annual IEEE International Conference on Communications (**ICC 2005**), May 2005, Seoul, Korea. To appear.
- **Dong Lu**, Huanyuan Sheng, Peter Dinda, "Size-Based Scheduling Policies with Inaccurate Scheduling Information", In Proceedings of 12th Annual Meeting of the IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (**MASCOTS 2004**), October 2004, Volendam, The Netherlands.
- Stefan Birrer, **Dong Lu**, Fabian Bustamante, Yi Qiao, Peter Dinda, "FatNemo: Building a Resilient Multi-Source Multicast Fat-Tree", In Proceedings of the Ninth International Workshop on Web Content Caching and Distribution (**WCW 2004**), Beijing, China. Also appeared in Springer LNCS.
- **Dong Lu**, P. Dinda "Synthesizing Realistic Computational Grids", in proceedings of 15th ACM/IEEE SuperComputing (**SC 2003**), Phoenix, AZ.
- **Dong Lu**, P. Dinda "GridG: Generating Realistic Computational Grids", ACM SIGMETRICS Performance Evaluation Review (**Per**), Volume 30, Number 4, 2003.
- **Dong Lu**, P. Dinda "Virtualized Audio: A Highly Adaptive Interactive High Performance Computing Application", In Proceedings of 6th Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers (**LCR 02**), Washington, DC, 2002. Also to appear in Springer LNCS.

### *Academic Activities*

- **Referee for Conferences and Workshops:** ACM/IEEE Supercomputing, IEEE HPDC, IEEE IWQoS, IEEE Cluster Computing, IEEE MASCOTS, IEEE ICDCS

### *Honors & Grants*

- Nominated by the Computer Science department for the **Presidential Fellowship** of Northwestern University, 2003.
- Global Grid Forum (GGF5) **student travel grant**, 2002.
- ACM Sigmetrics **student travel grant**, 2001.
- **Walter P. Murphy Fellowship**, Northwestern University, Evanston, IL, USA, 2000.
- **Scholarship of National Excellence**, Shanghai Jiao Tong (Chiao Tung) University, Shanghai, China, 1999.
- **Excellent Undergraduate Thesis**, Central South University, Changsha, China, 1997.
- **Scholarship of SME Educational Foundation**, Central South University, Changsha, China, 1997.
- **Torch Scholarship**, Central South University, Changsha, China, 1995.
- Six times **semester scholarship**, Central South University, Changsha, China, September 1993 ~ June 1997.
- Three times the title of **Excellent Student**, Central South University, Changsha, China, Sept. 1993 ~ June 1997.

Detailed Curriculum Vitae is available at: <http://www.cs.northwestern.edu/~donglu/CV/CV.html>

References available upon request.