

CS 525: Advanced Distributed Systems

(previously CS 598IG)

Course Overview

Over the past few decades, the functions of “traditional” operating systems have been scattered out to the edges of distributed systems. Peer-to-peer systems (think: Kazaa), sensor networks, the Grid, the Internet and the Web are examples. This course focuses on two case study areas: **peer-to-peer systems** and **sensor networks**. We will study efficient protocol design and evaluation, as well as learn high-level system issues. Research in these two areas also tends to be scattered across disjoint sets of researchers and conferences. A second emphasis of the course is on bridging the gap between these two focus areas and the existing base of **theoretical distributed computing**. This has the potential for applying algorithms or principles from one area to another; the course is an attempt to highlight these possibilities and motivate valuable projects.

Prerequisites

Basic Computer Science and basic computer programming skills are essential. Knowledge of Operating Systems (CS 241 or CS 423), or Networks (CS 438), or Distributed Systems (CS 425), or an equivalent course, or instructor consent, is required.

Course Website <http://www.cs.uiuc.edu/class/cs525>.

Timings

Class: Tuesday and Thursday, 9:30 AM - 10:45 AM, 1131 Siebel Center.

Office Hours (tentative, class days only): Tuesday and Thursday 10.45 AM - 12 PM, 3112 SC.

Course Staff

Dr. Indranil Gupta
3112 SC, indy@cs.uiuc.edu, 265-5517

Course Content

The first few weeks of lectures introduce ground basics in peer to peer systems, theory, and sensor network systems. Subsequently, about 80 research papers in various areas of distributed systems are presented, discussed, and debated by the students. The selection includes classical and contemporary papers from conferences including, but not limited to, PODC, SOSP, OSDI, Usenix, Infocom, SIGCOMM. Rather than running through all the papers in a few of such proceedings, we will pick and choose publications appropriate to the stated goals of this course.

An essential component of the course is a project involving at least one non-trivial idea and hands-on implementation. If this project leads to a conference-quality paper submission/acceptance, your course grade could benefit from it. You can collaborate in groups, and I will work with individual groups in defining the project and during its progress. At the end of the course, a top few “best” project papers will be selected, and given special attention for submission to conferences. For the record, 9 out of the 12 papers from the Fall 2003 course have been accepted in conferences (e.g., WCW, ICDCS, MASS and PODC), and many of the Fall 2004, Spring 2006, and Spring 2007

papers are either under conference review or have been accepted into conferences (e.g., Infocom, Middleware, QShine, CollaborateCom, GRID, etc.)!

About the Class

The initial few weeks of class will consist of lecturing, with the intent of building up common knowledge and grounding for the latter half of the course. We will then transition to student-led presentations of papers. Once student-led presentations start, students who are not presenting are expected to write short reviews (1-2 pages total) for any two of the “Main Papers” in that session. Active class participation is required, even in the initial part of the course!

Class Evaluation: Project, review papers, presentation, and class participation. Tentative splits are 50%, 20%, 25%, 5% respectively.

Abridged list of Topics (see course website for more comprehensive list): Epidemic algorithms, peer to peer systems, sensor networks, distributed computing theory, the Grid, overlays, routing, handling stress, distributed management, membership, classical algorithms, design methodologies, sources of unreliability and other characteristic studies, industrial systems, caching, publish-subscribe, structure of networks, selfish algorithms.