

Course Syllabus

 Edit

Welcome to OMS CS 6675 (summer 2024)

Instructor: Professor Dr. [Ling Liu](http://www.cc.gatech.edu/~lingliu/)  (<http://www.cc.gatech.edu/~lingliu/>)

Course Head TAs: Darren Mills

Course TAs: Thomas Lukose, Ryan (Yunxiang) Yan, Vijayraj Shanmugaraj

Course Objectives and Description

Advanced Internet-scale systems and applications are geographically distributed, highly available, incrementally scalable, and dynamically configurable. Typical questions that systems and advanced application developers are facing today include: How would you build a web service that can handle billions of frantic requests? What systems support do we need for developing applications of Internet scale? Can we provide dynamic configuration, replication, and migration of Web services? What new techniques will enable Internet systems and applications to better exploit high-speed networks? How should traditional systems issues such as naming, persistence, resource management, performance, and security be provided in a system of Internet scale? How much data can an internet scale system process? What does big data technology mean to a computer scientist? to a data scientist? to a business owner or a scientist?

This course reviews concepts, techniques, and systems issues in advanced Internet application development, and explores new challenges and research issues that are critical for developing Internet scale systems and applications. Main topics to be covered include fundamentals of search engines, (incl. robots and indexing servers), Web servers, application servers, web-based online transaction systems, content distribution networks, Internet scale social computing and social network systems, Blockchain systems, edge computing, mobile computing, Internet of Things, Digital twins. The coverage for each semester may vary. The course modules will provide the concrete list of topics covered by the current semester. One of the important goals of the course is to look beyond the present status of the Internet computing systems and applications, and conjecture new innovations for future Internet technologies and applications. The course will include a significant project component.

Prerequisites:

This class does not have mandatory prerequisites before participation. However, it covers broad categories of Internet computing systems technologies. Prior knowledge and work experiences in computer science, information technology, and applied engineering/science fields will give you additional

leverage towards learning and mastering broader and deeper concepts and methods covered in the course. See course [welcome \(https://gatech.instructure.com/courses/393412/pages/2022-spring-welcome-to-advanced-internet-computing-systems\)](https://gatech.instructure.com/courses/393412/pages/2022-spring-welcome-to-advanced-internet-computing-systems) page under before taking this class for more information.

Special note for summer session: This course is designed for regular session of 16 weeks. Summer session is 12 weeks total but we provide you the materials for all 26 modules. Due to the shortage of 4 weeks, the course project will be an optional instead of mandatory component in the summer session. We also adapt the grading policy to reflect this change for summer session.

Grading

Students will be evaluated by homework assignments, tests and project.

Grades will be computed using the tentative weighting scheme below:

- **Assignments: 50%**
- **Tests: 40%**
- **Participation: 10%**

NOTE: The deadline for all submissions (assignments and tests) will be 11:59pm Eastern on Sundays and we will have a graceful period with no penalty by 9am Eastern on Mondays to accommodate those in different time zones.

Written Assignments (50%)

The goal of the homework assignments is to help the students to master the knowledge learned from lectures through reading critiques on a selection of topics covered in the lectures. All homework assignments are individual assignments. They will be made available on Canvas under Assignment Tab in the beginning of the semester.

There are ten written assignments in the course: five Principles assignments and five Methods assignments. Each P assignment asks you to answer four questions, and all four questions are weighted equally. Each M assignment asks for a more thorough plan for performing research and development exercises on an Internet computing system/service/application from a list of given choices. These ten assignments together comprise 40% of your grade with each assignment worth 4% of your grade. All assignments should be written in MS word or google doc or Latex in single spacing, font size 12 pt and font size Arial or Times New Roman font face. All assignment submissions should be in PDF.

Tests (40%)

The goal of the tests is to consolidate the knowledge and learning experience of students about the

materials covered in the course modules. There are two proctored tests in this course. Given the summer semester is 11 weeks instead of 16 weeks, we will only cover 11 modules. Hence, the first covers the first half of the course (modules 1~5) and the second test covers the second half of the course (modules 6~11). Each question is multiple-choice, multiple-correct with five choices and between 1 and 4 correct answers. Partial credit is awarded. Each test is worth 20% of your overall grade. The tests are delivered via Canvas.

Participation (10%)

In this course, students will earn participation grade in several ways. The goal is to encourage each student in the class to interact with your peers, to participate in one another's assignment reviews and project check-in reviews, and to see the variety of approaches taken by your classmates in the assignments and projects. Examples of earning participation credits include completing peer reviews, participating in discussions on Ed Discussion. All types of participation are graded on not only quantity, but also quality. For example, peer reviews and Ed Discussion contributions only receive credit if they are substantive. Additional incentives are built in to encourage the completion of peer reviews quickly.

Projects (not required for Summer semester)

READ ME: During summer semester, we have 12 weeks instead of 16 weeks in the regular (fall or spring) semester, we place the project as an option. It is not a part of the grade for summer session, but if you may choose to do it in your spare time. If you would like to submit it for credit, you may inform the head GTA. We will consider giving you extra credits only when your project is reaching the quality level that deserves extra credits. Only showing the amount of efforts/time you put in does not translate to extra credits.

The rest of the project description is designed for regular semester where students have 3-4 weeks to work on their project.

There is one final project in this class. You can design your final project by choose a specific category of Internet computing technology from the materials covered from the course. Providing comprehensive investigation and redesign of an existing Internet system or application. Alternatively, you may also create a brandy new Internet service/application/system. You will conduct need finding by examining existing artifacts, like public forums, reviews, data logs, possibly combined with some research articles to identify the pros and cons of the existing system design in functional requirements (operational capabilities), non-functional requirements (e.g., security, privacy, trust, usability). You will also perform a heuristic evaluation of the existing application(s), grounded in the course's principles. You will then select one existing approach as the baseline design method. In the case of totally new project task, you will use the intuitive implementation as your baseline design method.

Based on the findings, you will complete a written document of your own prototype design, including the measurement of a working version of your own prototype (e.g., can be a revised version of the existing open source). Finally, you will provide a plan of how you would evaluate your prototype both qualitatively and empirically. The project is worth 20% of your final grade. The project should be written in MS word or google doc or latex in single spacing, font size of 12 pt, font face of Arial or Times New Roman. For more information, consults the project assignment or overview page on Canvas.

Useful References and Texts

There is no textbook required for this course. The course material comes primarily from course notes and a selection of recent papers on advanced Internet computing systems and application development and related research issues. However, there may be text books that you will find useful. A collection of papers and open access books will be provided at the course website from Canvas.

Useful Reference TextBooks (Not Required)

- [Building Secure and Reliable Network Applications](http://www.amazon.com/exec/obidos/ASIN/1884777295/qid=965867153/sr=1-1/102-7075309-9529766) , [Prof. Kenneth P. Birman](http://www.cs.cornell.edu/Info/Department/Annual95/Faculty/Birman.html) , <http://www.cs.cornell.edu/Info/Department/Annual95/Faculty/Birman.html>, Manning Pub. Co. 1996, Hardbound, 591 pages, ISBN 0137195842. It can be ordered from Follett DirectNet, 1-800-621-4088.
- [Understanding Search Engines: Mathematical Modeling and Text Retrieval \(Software, Environments, and Tools\)](http://www.amazon.com/exec/obidos/ASIN/0898714370/ref=cm_mp_wl/102-7075309-9529766?colid=3J8NNJ10YU5O4) , http://www.amazon.com/exec/obidos/ASIN/0898714370/ref=cm_mp_wl/102-7075309-9529766?colid=3J8NNJ10YU5O4, Michael W. Berry and Murray Browne.
- [Modern Information Retrieval](http://www.amazon.com/exec/obidos/ASIN/020139829X/o/qid=966447906/sr=2-1/102-7075309-9529766) , <http://www.amazon.com/exec/obidos/ASIN/020139829X/o/qid=966447906/sr=2-1/102-7075309-9529766>, Berthier Ribeiro-Neto, Ricardo Baeza-Yates (ACM Press Series) May 1999. Addison-Wesley Pub Co; ISBN: 020139829X.
- [The Developer's Guide to the Java\(tm\) Web Server\(tm\): Building Effective and Scalable SeverSide Applications](http://www.amazon.com/exec/obidos/ASIN/020137949X/o/qid=966441348/sr=2-2/102-7075309-9529766) , <http://www.amazon.com/exec/obidos/ASIN/020137949X/o/qid=966441348/sr=2-2/102-7075309-9529766>, Dan Woods, Larne Pekowsky, Tom Snee, and Connie Welss
- [Web Server Technology](http://www.amazon.com/exec/obidos/ASIN/155860376X/102-7075309-9529766) , <http://www.amazon.com/exec/obidos/ASIN/155860376X/102-7075309-9529766>, Nancy J. Yeager and Robert E. McGrath, Morgan Kaufmann Publishers, 1996.
- [Web Proxy Servers](http://www.amazon.com/exec/obidos/ASIN/0136806120/102-7075309-9529766) , <http://www.amazon.com/exec/obidos/ASIN/0136806120/102-7075309-9529766>, Ari Luotonen, Prentice Hall, 1998
- [The Grid: Blueprint for a New Computing Infrastructure](http://www.amazon.com/exec/obidos/ASIN/1558604758/102-7075309-9529766) , <http://www.amazon.com/exec/obidos/ASIN/1558604758/102-7075309-9529766>, Edited by Ian Foster and Carl Kesselman, Morgan Kaufmann Publishers, 1998.

Useful and Related Links

Writing and Presentation Style:

- [The Elements of Style](http://www.bartleby.com/141/) ↗ (http://www.bartleby.com/141/)
- [Advice on Research and Writing](http://www.cs.cmu.edu/afs/cs.cmu.edu/user/mleone/web/how-to.html) ↗ (http://www.cs.cmu.edu/afs/cs.cmu.edu/user/mleone/web/how-to.html)
- [On Being a Scientist](http://www.nap.edu/readingroom/books/obas) ↗ (http://www.nap.edu/readingroom/books/obas)
- [How to write a great research paper](https://www.microsoft.com/en-us/research/academic-program/write-great-research-paper/) ↗ (https://www.microsoft.com/en-us/research/academic-program/write-great-research-paper/) by Simon Peyton Jones
- [How to Increase the Chances Your Paper is Accepted at ACM SIGCOM](http://ccr.sigcomm.org/archive/1998/jul98/ccr-9807-partridge.pdf) ↗ (http://ccr.sigcomm.org/archive/1998/jul98/ccr-9807-partridge.pdf) by Craig Partridge

Internet Resources and Tools

- [Internet Traffic Report \(World Wide\)](http://www.internettrafficreport.com/) ↗ (http://www.internettrafficreport.com/)
- [The Internet Archive](http://www.archive.org/) ↗ (http://www.archive.org/)
- [The World Wide Web History Project](http://www.webhistory.org/home.html) ↗ (http://www.webhistory.org/home.html)
- [Z39.50 and the World Wide Web](http://www.dlib.org/dlib/march96/briefings/03indexdata.html) ↗ (http://www.dlib.org/dlib/march96/briefings/03indexdata.html)
- [An Internet Encyclopedia \(Third Edition\)](http://www.freesoft.org/CIE/index.htm) ↗ (http://www.freesoft.org/CIE/index.htm)
- [Internet Statistics Growth and Usage of the Web and the Internet](http://www.mit.edu/people/mkgray/net/) ↗ (http://www.mit.edu/people/mkgray/net/)
- [Search Engine Watch](http://www.searchenginewatch.com/) ↗ (http://www.searchenginewatch.com/)
- [Computing Research Repository \(CoRR\)](http://arxiv.org/archive/cs/intro.html) ↗ (http://arxiv.org/archive/cs/intro.html)