

Fall 2018 – CSCI 8000 (40997)

Advanced Cloud Computing

In Kee Kim
Department of Computer Science
The University of Georgia
inkee.kim@uga.edu

August 13, 2018

1 General Information

- Instructor: In Kee Kim (inkee.kim@uga.edu)
- Office Location: 549A @Boyd GSRC
- Office Hours: Thursday 3:30 p.m. – 5:00 p.m. (after class)
- Credits: 4
- Course website: <http://cobweb.cs.uga.edu/~kim/classes/F18-CSCI8000/>
- Class meeting time and location:
 - Tue and Thurs: 2:00 – 3:15 p.m. @Boyd GSRC 0208
 - Wed: 2:30 – 3:20 p.m. @Boyd GSRC 0208

2 Course Overview

Cloud Computing has become a *de facto* infrastructure in many business and research organizations to deliver various user-facing, business, and scientific applications to end users. The goal of this course is to introduce the underlying technologies that created the current cloud computing and infrastructure, then discuss the future of clouds by investigating the *cutting-edge* research that will soon be part of cloud ecosystems.

This is a graduate-level, research-oriented, and a (mostly) seminar-type course in cloud computing. The instructor will introduce the fundamentals of cloud computing in the first a few lectures. Then students will read recent papers and discuss the following topics. i.e., modern cloud infrastructures and applications, cloud-scale machine learning, IoT/Edge, and the future of clouds.

In addition, the students will have opportunities to present assigned research papers. The students' presentation should clearly address and summarize motivation, problem statement, the authors' approach, evaluation, and discussion topics (e.g., pros and cons, further improvement, etc.).

Prerequisite: Not required, but prior knowledge of operating systems, distributed systems, and computer networks will be a plus.

Textbooks: This class does not require a textbook. Course materials will be recent publications from top-tier cloud/system conferences and journals.

3 Course Topics

In this course, we will study a variety of basic principles, techniques, and modern advances in cloud computing, large-scale distributed systems, (cloud-scale) machine learning systems, and IoT/Edge computing. Topics to be covered include (schedule is *tentative* and subject to change, please keep track of it on the course website):

1. Course overview and introduction to cloud computing (~1 week)
2. Cloud 101: Fundamentals (~3 weeks)
 - Cloud service models
 - Data centers, cloud infrastructure, VM, and containers
 - Elasticity and autoscaling
 - Cloud applications and workloads
3. Cloud infrastructure (~4 weeks)
 - Large-scale cluster management. i.e., resource sharing, scheduling, provisioning
 - Container orchestration and microservice management
 - Serverless computing and cloud functions
4. Cloud Computing and ML/AI (~4 weeks)
 - Large-scale machine learning service on clouds
 - Resource management for production-scale machine learning
 - Applied machine learning research in clouds. i.e., workload and error prediction
 - Current and future (system specific) research challenges for AI and ML
5. Cloud IoT and Edge (~2 weeks)
 - Cloud IoT and edge computing fundamentals
 - Edge computing applications
 - Future research direction/opportunity in the cloud and edge computing

4 Class Presentation

The presentation should be prepared for 30–40 minutes talk and clearly explain motivation, problem statement, related work (additional literature survey is a plus), the authors’ approach, evaluation, and discussion topics (e.g., pros and cons, further improvement, etc.).

Slide for the presentation must be emailed to the instructor by 11 p.m. before the class day. Also, students can always consult with the instructor for the presentation preparation.

5 Assignments and Final Research Project

Reading Assignments (Paper Review): All students should read each assigned paper and submit the review before the scheduled lecture (9 a.m. of the day). Please note that the reading assignment will be started in the second week of the semester. The paper review should have three paragraphs:

- **Paragraph 1:** 1–2 sentences of problem statement and 1–2 summary of the paper
 1. e.g., **Problem Statement:** what problem is being addressed in the paper?
 2. e.g., **Summary of Approach/Evaluation:** “the authors address the problem by designing and implementing...”
- **Paragraph 2:** pros and cons of the paper
 1. e.g., **Pros:** what’s compelling/effective about the paper.
 2. e.g., **Cons:** what are the most significant flaws/limitations of the paper?
- **Paragraph 3:** (at least) three questions for the paper discussion

Final Research Project: An essential component of this course is a final research project. You are expected to work on a cutting-edge research problem, which is relevant to the topic of this course. You can work either individually or with other students (1 to 3 members). Project ideas have to be discussed with the instructor. Tentative schedule for this project is as follows:

- **Week 5:** Introduction to final project (including sample projects) by the instructor
- **Week 7:** Project proposal due
- **Week 8:** Project proposal presentation
- **Week 12:** Mid-phase presentation or meeting
- **Week 15 (a week after Thanksgiving break):** Final report and presentation

6 Grading Distribution

- Class participation: 10%
- Paper presentation: 20%
- Reading assignment: 20%
- Two project reports: 30% (10% for proposal, 20% for final report)
- Final project presentation: 20%

7 Academic Honesty

All students must follow the Academic Honesty Policy of the University of Georgia. The detailed information of this policy can be found at <https://honesty.uga.edu/Academic-Honesty-Policy/>. If there are any issues regarding this policy, please contact the instructor immediately.