Course Syllabus

Edit

CS 643 1J2 Spring 2024 - Syllabus

Instructor

- Cristian Borcea
- borcea@njit.edu (mailto:borcea@njit.edu)
- https://cs.njit.edu/~borcea (https://cs.njit.edu/~borcea)
- Wednesdays, 6pm-8:50pm, Jersey City Campus.
- Office hours: before and after class. Also, weekly online: Mondays, 5:15pm-6pm, Webex: https://njit.webex.com/meet/borcea (https://njit.webex.com/meet/borcea)
  - You can also join by phone: 1-650-479-3207 *Call-in toll number (US/Canada).* Access code: 928 115 913

Short Description

The course presents a comprehensive view of cloud computing, from platforms and services to programming and infrastructure. The topics include: cloud computing platforms, with examples from Amazon Web Services (AWS), Google Cloud, and Microsoft Azure; cloud services for data analytics, machine learning, mobile computing, IoT, edge computing, security and privacy, and devops; programming frameworks for parallel computing in the cloud; distributed storage in the cloud; and virtualization and containerization. The course includes homework, programming assignments, and research paper presentations. The programming assignments will be done in AWS.

Learning Outcomes

Upon the successful completion of this course, the student should be able to:

- Analyze the trade-offs between deploying applications/services in the cloud and on premises
- Compare the advantages and disadvantages of different types of cloud platforms
- Deploy applications in public cloud platforms
- Program data intensive parallel applications in the cloud
- Analyze the performance, scalability, and availability of cloud systems and applications
- Identify security and privacy issues in the cloud
- Present state-of-the-art cloud research

Why Take This Course?

Cloud computing represents a major paradigm shift in computing from the era of personal computers to the era of computing as utility. Most major Internet services are already deployed in the "the cloud."

We store most of our data in "the cloud" and execute most applications from "the cloud." This course is aimed at all graduate students (both M.S. and Ph.D. students) who want to learn how to design and program cloud services and applications as well as how to build and administer cloud systems. By studying real-world systems developed in industry, students will acquire cutting-edge knowledge that may be a major advantage when searching for a job.

Lectures and Readings

There is no book required for this class. Each lecture is based on research papers and/or online documentation covering a specific topic (i.e., readings). The slides for each lecture will be posted before each class. After each class, the voice-over video recordings of each sub-module of the lecture will be posted as well. The readings should be used as reference material to clarify and add details to lectures.

Exams

There will be two in-class exams: a midterm, and a final exam. Both exams are closed book (i.e., papers, notes). The final exam will cover only the material taught after the midterm. In case of missing an exam, a make-up may be taken only after providing written documentation to the Dean of Students.

Homework

Homework will be assigned 6 times during the semester to prepare students with the type of questions encountered in exams. The solutions will be discussed in class.

Programming Assignments

There will be two individual programming assignments. The first is to build an image recognition pipeline in Amazon AWS, using two EC2 instances, S3, SQS, and Rekognition. The assignment must be done in Java on Amazon Linux VMs. You will learn how to use the AWS cloud platform and how to develop an AWS application that uses existing cloud services. The second is to build a machine learning prediction model in Spark/MLlib over AWS. The model must be trained in parallel on multiple EC2 instances. The assignment must be implemented in Java, Scala, or Python on Ubuntu Linux. You will learn how to develop parallel machine learning applications in the AWS cloud platform.

Research Paper Presentations

Students will present individually one research paper during the semester. These papers cover state-of-the-art research in cloud computing. The presentations will be in class in week 7, and the slides must be uploaded in Canvas before the presentations. Extra-credit is available for asking good questions about the presentations.

Grading

- Midterm exam - 25%
- Final exam - 25%
- Programming Assignment 1 - 10%
- Programming Assignment 2 - 20%
- Research presentation - 8%
- Homework - 12%

Schedule

- **Week 1: Course overview. Introduction to Cloud Computing.**
- **Week 2: Cloud Platforms I: Infrastructure as a Service (IaaS), AWS.**
  - Homework 1 handed out on 25 January
  - Research papers assigned by 25 January
- **Week 3: Cloud Platforms II: Serverless Computing; Function as a Service (FaaS).**
  - Homework 1 due on 1 February
  - Programming assignment 1 handed out on 1 February
- **Week 4: Cloud Services I: Data Analytics and Machine Learning.**
  - Homework 2 handed out on 8 February
- **Week 5: Cloud Services II: Mobile, IoT, and Edge Computing.**
  - Homework 2 due on 15 February
- **Week 6: Cloud Services III: Security and Privacy, Devops.**
  - Homework 3 handed out on 22 February
  - Programming assignment 1 due on 22 February
- **Week 7: Research Paper Presentations**
  - Homework 3 due on 29 February
- **Week 8: Midterm on 7 March**
  - Discussion of midterm solutions in class after the midterm
- **Week 9: Parallel Programming in the Cloud I: Google's MapReduce, Apache's Hadoop, Yahoo's Pig Latin.**
  - Programming assignment 2 handed out on 21 March
- **Week 10: Parallel Programming in the Cloud II: Apache's Spark, Storm and Zookeper.**
  - Homework 4 handed out on 28 March
- **Week 11: Cloud Storage Systems I: Google's GFS and BigTable.**
  - Homework 4 due on 4 April
- **Week 12: Cloud Storage Systems II: Amazon's Dynamo and Other Cloud Databases.**
  - Homework 5 handed on 11 April
- **Week 13: Virtualization I: VMWare, XEN, Live VM Migration.**
  - Homework 5 due on 18 April
- **Week 14: Virtualization II: Containerization, Docker, Kubernetes.**
  - Homework 6 handed out on 25 April and due on 2 May

- Programming assignment 2 due on 25 April
- **Week 15: Final Exam is on 9 May**

Academic Integrity

Academic Integrity is the cornerstone of higher education and is central to the ideals of this course and the university. Cheating is strictly prohibited and devalues the degree that you are working on. As a member of the NJIT community, it is your responsibility to protect your educational investment by knowing and following the academic code of integrity policy that is found at this link: **University Policy on Academic Integrity (https://www5.njit.edu/policies/sites/policies/files/academic-integrity-code.pdf)** .

Please note that it is my professional obligation and responsibility to report any academic misconduct to the Dean of Students Office. Any student found in violation of the code by cheating, plagiarizing or using any online software inappropriately will result in disciplinary action. This may include a failing grade of F, and/or suspension or dismissal from the university. If you have any questions about the code of Academic Integrity, please contact the Dean of Students Office at **dos@njit.edu. (mailto:dos@njit.edu.)**

Modifications to Syllabus

The students will be consulted and must agree to any modifications or deviations from the syllabus throughout the course of the semester.