

Cloud Computing

2024/2



Class time: To be defined

Time commitment: 8 hours/week

Academic credit: 4

Course format: Lecture

Asynchronous study: Students who cannot attend the lectures and assignment sessions synchronously may watch the recorded videos. Our lectures are expected to have in-depth discussions on the day's subject while the pure information-based lecture content will be available as video recordings.

Instructor's information

Ítalo Cunha

Associate professor, Universidade Federal de Minas Gerais, Brazil

E-mail: cunha@dcc.ufmg.br

Webpage: <http://www.dcc.ufmg.br/~cunha>

Ítalo Cunha is an associate professor at the Computer Science Department at UFMG, Brazil. He developed his Ph.D. research at Technicolor Research and Innovation and graduated from Université Pierre et Marie Curie in 2011. Ítalo worked as an Associate Visiting Research Scholar at Columbia University during 2019. His research focuses on improving network performance, reliability, and security. His contributions provide better visibility on Internet topology and routing dynamics; help network operators troubleshoot failures and performance problems; support identification and mitigation of cybersecurity threats; and empower novel networking and security research.

What is this course about?

Cloud computing is the delivery of computing resources or applications over a network, as a utility (like water or electricity). Cloud computing allows customers and enterprises to pay for resources they use, while not owning any infrastructure. Cloud computing has revolutionized how applications are delivered to users and enterprises; reducing costs, improving availability, allowing dynamic scaling on demand, and improving performance.

In this course we will discuss the motivations and trade-offs of moving to the cloud. We will study the technologies behind the main components of cloud computing: infrastructure, resource sharing, development platforms, and applications. The cloud computing *infrastructure* is based on large datacenters that benefit from economies of scale; *resource sharing* is achieved through efficient virtualization of compute, network, and storage; flexible and powerful frameworks provide *development platforms* supporting different paradigms for building *applications*. This course will cover these pieces in depth and how they combine to support cloud computing. Students will also obtain hands-on experience with a set of programming assignments that will employ modern cloud computing technologies and exercise best programming/DevOps practices.

What background knowledge do I need before taking this course?

Prerequisites: DCC205 (Data Structures) or an equivalent course.

You are supposed to have experience in a high-level programming language and in employing data-structures, as taught in DCC205. Understanding of computer architecture, computer networks, and databases as taught in DCC006 (Computer Architecture 1), DCC023 (Computer Networks) and DCC011 (Introduction to Databases), are beneficial but not assumed.

What will I learn in this course?

A student who has met the objectives of this course will be able to:

1. Explain what cloud computing is; discuss its benefits and limitations regarding power, efficiency, scale, cost, and performance; and evaluate when migration to the cloud is beneficial (compared to local compute) for a given application or use case.
2. Describe, evaluate and compare the different solutions available in cloud computing, including virtualization of CPU, memory, storage, and the network; monitoring, management, automation solutions; and programming paradigms such as MapReduce, microservices, and serverless computing.
3. Compare the advantages and disadvantages of different cloud computing platforms.
4. Design, develop, and manage dynamically scaled, parallel, and distributed cloud native applications. Consider scalability, costs, and performance constraints. Identify security and privacy issues in cloud applications.
5. Use cloud computing to perform big data analytics using different APIs, frameworks, paradigms, and libraries.

What will I do in this course?

Mastery of the subject matter of this course will be evaluated in the following ways:

Programming and DevOps assignments: Students will develop and deploy services to exercise course material, in particular: monitoring, elasticity, and management of a cloud native application; use and compare different storage approaches; and develop data analytics applications in multiple programming frameworks.

Short video on advanced topics related to cloud computing: Students will form groups to review a recent research paper related to cloud computing and prepare a short video describing the technology and relating it to material seen in the course.

Midterm exam: There will be a midterm exam in this class covering topics discussed in the first 3 weeks. The midterm will focus on the conceptual and theoretical aspects of the course. After each exam you will have the opportunity to boost your grade by designing one new challenging exam question and answer.

Comprehensive final exam: A comprehensive final exam will be designed to test students' overall understanding of the whole course materials, with a focus on the theoretical aspects of the course. It is designed to be a comprehensive exam with a longer exam time. It covers everything you have learned from this class.

How can I prepare for the class sessions to be successful?

- Prior reading of the textbook materials before class is always highly recommended. Engaging in discussions during class is encouraged.
- It is your responsibility to check the course website for any important announcements. You will need to download course slides, assignments, instructions, and any other necessary information from Moodle.

- It is recommended that you should start to work on every assigned course work early. Allocate time to work on assignments each week to make continued progress.
- After the midterm exam grade is announced, you are encouraged to craft one new challenging exam question and answer as an optional bonus (5 bonus points) to boost your midterm grade.

What required texts, materials, and equipment will I need?

The main book for this class is:

[DC] The Cloud Computing Book: The Future of Computing Explained. Douglas Comer. Chapman and Hall/CRC, 1st ed.

Additional required reading material and resources will be provided on the course website.

What optional texts or resources might be helpful?

Additional supplementary material will be added to Sakai. Other recommended open-access material:

[FG] Cloud Computing for Science and Engineering. Ian Foster and Dennis B. Gannon. MIT Press, 1st ed. Available at: <https://cloud4scieng.org/chapters/>

[DWDL] Learning Spark. Jules S. Damji, Brooke Wenig, Tathagata Das, and Denny Lee. O'Reilly, 2nd ed. Available at: <https://databricks.com/p/ebook/learning-spark-from-oreilly>

[MMDS] Mining of Massive Datasets. Jure Leskovec, Anand Rajaraman, and Jeff Ullman. Cambridge University Press, 3rd ed. Available at: <http://www.mmds.org/>

[OPS] Edge Cloud Operations: A Systems Approach. Peterson, Baker, Bavier, Williams and Davie. <https://ops.systemsapproach.org/>

Additional public resources will be used in the programming assignments and practice exercises.

How will my grade be determined?

Total grade is 100%, composed of:

- Programming assignments: 12.5% points each, for a total of 50%
- Midterm exam: 15%
- Final exam: 25%
- Video: 10%

What are the course policies?

Communications:

Email is the best way to contact me (cunha@dcc.ufmg.br). In general, you should expect a response to your message within 48 hours during weekdays or 72 hours on the weekends. Please include [DCC-CloudComp] in the email subject line. Another way to get help outside lectures is to post questions on the online forums/discussions.

Assignment Deadlines:

Late coursework submissions will not be accepted unless you communicate with me prior to the due date and justify the need for an extension.

Discussion Guidelines:

Civility is an essential ingredient for academic discourse. All communications for this course should be conducted constructively, civilly, and respectfully. Differences in beliefs, opinions, and approaches are to be expected. Please bring any communications you believe to be in violation of this class policy to the attention of your instructor. Active interaction with peers and your instructor is essential to success in this course, paying particular attention to the following:

- Be respectful of others and their opinions, valuing diversity in backgrounds, abilities, and experiences.
- Challenging the ideas held by others is an integral aspect of critical thinking and the academic process. Please word your responses carefully and recognize that others are expected to challenge your ideas. A positive atmosphere of healthy debate is encouraged.
- Read your online discussion posts carefully before submitting them.

Academic Integrity:

As a student, you should abide by the academic honesty standard of the university. For all graded work, students should pledge that they have neither given nor received any unacknowledged aid.

Estimated course schedule?

Date	Class topic/unit name	Pre-class work for students	Planned in-class activities	Assignments due
Weeks 1-2	Foundations, motivation, use cases, services, infrastructure	[DC] Chapters 1-4	Introduction Lecture	Assignment 1 announced
Week 3	Virtualization I, MapReduce	[DC] Chapters 5, 11	Lecture	Assignment 1 due, Assignment 2 announced
Week 4	Automation and Orchestration	[DC] Chapters 9, 10	Lecture	
Weeks 5-6	Programming paradigms	[DC] Chapters 12, 13, 14	Lecture Midterm	Assignment 2 due, Assignment 3 announced
Weeks 7-8	Virtualization II	Chapters 6-8	Lecture	Assignment 3 due, assignment 4 announced

Week 9-10	Monitoring, orchestration, and DevOps	[DC] Chapters 10, 15	Lecture	Video assignment announced
Week 11-12	Edge computing, security, privacy, complexity	[DC] Chapters 16, 17, 18; [OPS] Chapters 3-6	Lecture	Assignment 4 due, video due