

Course Name : Distributed Systems and Cloud Computing							
Course Code	Course Type	Regular Semester	Lecture (hours/week)	Seminar (hours/week)	Lab. (hours/week)	Credits	ECTS
CMP 408	B	Spring	3.00	1.00	0.00	3.50	6.00
Lecturer Artur Koci, PhD							
Assistant Hersi Kopani, Msc							
Course language Albanian							
Course level Master							
Description This course covers the systems architecture and enabling technologies of parallel and distributed computing systems and their innovative applications. We cover scalable multiprocessor systems, clusters of computers, P2P networks, computational Grids, virtual machines, virtual clusters, Internet Clouds, Internet of Things (IoT) and social networks. Case studies include Sequoia, Tianhe-1A, TeraGrids, XEN, Google AppEngine, Amazon EC2 and S3, NASA and CERN clouds, BitTorrent, Facebook, Twitter, MapReduce, Hadoop, VMWare Tools, and some cloud and IoT projects, etc.							
Objectives At the end of the course, the student will be able to: -Understand the main concepts and models of distributed systems and cloud computing. -Apply knowledge in the architecture of parallel and distributed systems. -Analyze virtualization, virtual classrooms and virtualized datacenters. -Implement cloud-based solutions, such as AWS, Google Cloud and Microsoft Azure. -Evaluate security, resource management and performance techniques in cloud and distributed systems.							
Core Concepts Distributed Architectures and Models Communication, Coordination, and Consistency Resource Management and Virtualization Security and Reliability Data Storage and Processing							
Course Outline							
Week	Topic						
1	This week covers Scalable Computing over the Internet, Technologies for Network-based Systems, and System Models for Distributed and Cloud Computing. It explores the fundamental concepts of distributed systems, enabling technologies, and key models for distributed and cloud computing. Study: Chapter 1 (Pages 4-36)						
2	This week covers Software Environments for Distributed Systems and Clouds and Performance, Security, and Energy Efficiency. It explores the role of software platforms in distributed and cloud systems, along with key aspects of system performance, security challenges, and energy-efficient computing. Study: Chapter 1 (Pages 36-65)						
3	This week covers Clustering for Massive Parallelism and Computer Clusters and MPP Architectures. It explores the fundamentals of clustering techniques used to achieve large-scale parallel computing and the architectural design of massively parallel processing (MPP) systems. Study: Chapter 2 (Pages 65-87)						
4	This week covers Design Principles of Computer Clusters, Cluster Job and Resource Management, and Case Studies of Top Supercomputer Systems. It focuses on the fundamental principles of cluster computing, job scheduling, resource allocation, and real-world implementations of supercomputing clusters. Study: Chapter 2 (Pages 87-122)						
5	This week covers Implementation Levels of Virtualization and Virtualization Structures, Tools, and Mechanisms. It introduces key concepts of virtualization, different levels of implementation, and essential tools used for virtualized environments in distributed computing. Study: Chapter 3 (Pages 130-155)						

6	This week covers Virtualization of CPU, Memory, and I/O Devices, Virtual Clusters and Resource Management, and Virtualization for Data-Center Automation. It focuses on the impact of virtualization on computing resources, efficient management of virtual clusters, and automation techniques in cloud data centers. Study: Chapter 3 (Pages 155-183)
7	This week covers Cloud Platform Architecture over Virtualized Data Centers. The discussion includes cloud computing concepts and service models, data center design and interconnection networks, and the architectural design of compute and storage clouds. Strategies for resource optimization and the scalability of cloud platforms are analyzed. Study: Chapter 4 (Pages 192-215)
8	Midterm
9	This week covers Public Cloud Platforms, Inter-cloud Resource Management, and Cloud Security and Trust Management. The discussion includes major public cloud platforms such as Google App Engine (GAE), Amazon Web Services (AWS), and Microsoft Azure. Additionally, inter-cloud resource management techniques are explored, focusing on workload distribution and interoperability. The week concludes with an analysis of cloud security challenges and trust management strategies. Study: Chapter 4 (Pages 227-265)
10	This week covers Services and Service-Oriented Architecture, Message-Oriented Middleware, and Portals and Science Gateways. The discussion includes the concept of service-oriented architecture (SOA), the role of message-based middleware in data exchange, and the use of portals and science gateways for resource access and management in distributed systems. Study: Chapter 5 (Pages 272-304)
11	This week covers Discovery, Registries, Metadata, and Databases, along with Workflows in Service-Oriented Architectures. The discussion focuses on the importance of registries and metadata in service management, as well as the use of workflow models for coordinating and executing distributed services. Study: Chapter 5 (Pages 304-331)
12	This week covers Cloud Programming and Software Environments. The discussion includes the features of cloud and grid platforms, parallel and distributed programming paradigms, and programming support for Google App Engine, Amazon AWS, and Microsoft Azure. Additionally, emerging cloud software environments and development trends are analyzed. Study: Chapter 6 (Pages 336-405)
13	Introduction to AWS and Core Services This week focuses on getting familiar with Amazon Web Services (AWS) and its core services. Topics include AWS global infrastructure, identity and access management (IAM), compute services like EC2 and Lambda, and data storage solutions such as Amazon S3 and EBS. Networking concepts, including Amazon VPC, Route 53, and load balancing, are also discussed. Hands-on activities include creating an AWS Free Tier account, launching an EC2 instance, configuring IAM policies, and storing data in Amazon S3.
14	This week shifts the focus to managing and deploying applications on AWS. Topics include AWS database services such as RDS, DynamoDB, and Redshift, as well as serverless computing with AWS Lambda and API Gateway. Cloud monitoring and security tools, including CloudWatch, CloudTrail, and AWS Security Hub, are covered, along with DevOps automation for CI/CD using AWS CodePipeline and CodeDeploy. Hands-on activities involve deploying a serverless application, setting up an RDS database, and implementing basic security monitoring in AWS.
15	Presentation of the project
16	Final Exam

Prerequisites	The student must attend the course at a minimum rate of 75%.
Literature	<ul style="list-style-type: none"> Hwang, Kai Dongarra, Jack Fox, Geoffrey C- Distributed and cloud computing: from parallel processing to the internet of things
References	<ul style="list-style-type: none"> Kai Hwang, Min Chen - Big-Data Analytics for Cloud, IoT and Cognitive Computing https://www.aws.training/

Course Outcome

1	Understand and Apply Distributed and Cloud Architectures
2	Develop and Deploy Cloud-Based Applications
3	Manage Distributed Computing Resources Efficiently
4	Ensure Security and Reliability in Distributed Systems
5	Analyze and Optimize Data Processing in Distributed and Cloud Environments

Course Evaluation

In-term Studies	Quantity	Percentage
Midterms	1	20
Quizzes	0	0
Projects	0	0
Term Projects	1	30
Laboratory	0	0
Class Participation	0	0
Total in-term evaluation percent		50
Final exam percent		50
Total		100

ECTS Workload (Based on Student Workload)

Activities	Quantity	Duration (hours)	Total (hours)
Course duration (Including the exam week: 16x Total hours of the course)	16	4	64
Study hours outside the classroom (Preparation, Practice, etc.)	14	3	42
Duties	1	10	10
Midterms	1	15	15
Final Exam	1	19	19
Other	0	0	0
Total Work Load			150
Total Work Load / 25 (hours)			6.00
ECTS			6.00