



1. COURSE

CS3P2. Cloud Computing (Mandatory)

2. GENERAL INFORMATION

2.1 Course	:	CS3P2. Cloud Computing
2.2 Semester	:	10 ^{mo} Semestre.
2.3 Credits	:	3
2.4 Horas	:	1 HT; 4 HP;
2.5 Duration of the period	:	16 weeks
2.6 Type of course	:	Mandatory
2.7 Learning modality	:	Blended
2.8 Prerequisites	:	CS370. Big Data. (9 th Sem) CS370. Big Data. (9 th Sem)

3. PROFESSORS

Meetings after coordination with the professor

4. INTRODUCTION TO THE COURSE

In order to understand the advanced computational techniques, the students must have a strong knowledge of the various discrete structures, structures that will be implemented and used in the laboratory in the programming language.

5. GOALS

- That the student is able to model computer science problems using graphs and trees related to data structures.
- That the student apply efficient travel strategies to be able to search data in an optimal way.

6. COMPETENCES

- 1) Analyze a complex computing problem and to apply principles of computing and other relevant disciplines to identify solutions. (**Usage**)
- 6) Apply computer science theory and software development fundamentals to produce computing-based solutions. (**Usage**)

7. TOPICS

Unit 1: Distributed Systems (15)	
Competences Expected:	
Topics	Learning Outcomes
<ul style="list-style-type: none"> • Faults (cross-reference OS/Fault Tolerance) <ul style="list-style-type: none"> – Network-based (including partitions) and node-based failures – Impact on system-wide guarantees (e.g., availability) • Distributed message sending <ul style="list-style-type: none"> – Data conversion and transmission – Sockets – Message sequencing – Buffering, retrying, and dropping messages • Distributed system design tradeoffs <ul style="list-style-type: none"> – Latency versus throughput – Consistency, availability, partition tolerance • Distributed service design <ul style="list-style-type: none"> – Stateful versus stateless protocols and services – Session (connection-based) designs – Reactive (IO-triggered) and multithreaded designs • Core distributed algorithms <ul style="list-style-type: none"> – Election, discovery 	<ul style="list-style-type: none"> • Distinguish network faults from other kinds of failures [Familiarity] • Explain why synchronization constructs such as simple locks are not useful in the presence of distributed faults [Familiarity] • Write a program that performs any required marshalling and conversion into message units, such as packets, to communicate interesting data between two hosts [Usage] • Measure the observed throughput and response latency across hosts in a given network [Usage] • Explain why no distributed system can be simultaneously consistent, available, and partition tolerant [Familiarity] • Implement a simple server – for example, a spell checking service [Usage] • Explain the tradeoffs among overhead, scalability, and fault tolerance when choosing a stateful v stateless design for a given service [Familiarity] • Describe the scalability challenges associated with a service growing to accommodate many clients, as well as those associated with a service only transiently having many clients [Familiarity] • Give examples of problems for which consensus algorithms such as leader election are required [Usage]
Readings : [Cou+11]	

Unit 2: Cloud Computing (15)	
Competences Expected:	
Topics	Learning Outcomes
<ul style="list-style-type: none"> • Visión global de <i>Cloud Computing</i>. • Historia. • Visión global de las tecnologías que envuelve. • Beneficios, riesgos y aspectos económicos. • Cloud services <ul style="list-style-type: none"> – Infrastructure as a service <ul style="list-style-type: none"> * Elasticity of resources * Platform APIs – Software as a service – Security – Cost management • Internet-Scale computing <ul style="list-style-type: none"> – Task partitioning – Data access – Clusters, grids, and meshes 	<ul style="list-style-type: none"> • Explicar el concepto de Cloud Computing. [Familiarity] • Listar algunas tecnologías relacionadas con Cloud Computing. [Familiarity] • Explain strategies to synchronize a common view of shared data across a collection of devices [Familiarity] • Discutir las ventajas y desventajas del paradigma de Cloud Computing. [Familiarity] • Expresar los beneficios económicos así como las características y riesgos del paradigma de Cloud para negocios y proveedores de cloud. [Familiarity] • Diferenciar entre los modelos de servicio. [Usage]
Readings : [HDF11], [BVS13]	

Unit 3: Centros de Procesamiento de Datos (10)	
Competences Expected:	
Topics	Learning Outcomes
<ul style="list-style-type: none"> • Visión global de un centro de procesamiento de datos. • Consideraciones en el diseño. • Comparación de actuales grandes centros de procesamiento de datos. 	<ul style="list-style-type: none"> • Describir la evolución de los Data Centers. [Familiarity] • Esbozar la arquitectura de un data center en detalle. [Familiarity] • Indicar consideraciones de diseño y discutir su impacto. [Familiarity]
Readings : [HDF11], [BVS13]	

Unit 4: Cloud Computing (20)	
Competences Expected:	
Topics	Learning Outcomes
<ul style="list-style-type: none"> • Virtualization <ul style="list-style-type: none"> – Shared resource management – Migration of processes • Seguridad, recursos y aislamiento de fallas. • Almacenamiento como servicio. • Elasticidad. • Xen y Wmware. • Amazon EC2. 	<ul style="list-style-type: none"> • Virtualization <ul style="list-style-type: none"> – Shared resource management – Migration of processes . [Familiarity] • Explain the advantages and disadvantages of using virtualized infrastructure. [Familiarity] • Identificar las razones por qué la virtualización está llegando a ser enorme útil, especialmente en la cloud. [Familiarity] • Explicar diferentes tipos de aislamiento como falla, recursos y seguridad proporcionados por la virtualización y utilizado por la cloud. [Familiarity] • Explicar la complejidad que puede tener el administrar en términos de niveles de abstracción y interfaces bien definidas y su aplicabilidad para la virtualización en la cloud. [Familiarity] • Definir virtualización y identificar diferentes tipos de máquinas virtuales. [Familiarity] • Identificar condiciones de virtualización de CPU, reconocer la diferencia entre <i>full virtualization</i> y <i>paravirtualization</i>, explicar emulación como mayor técnica para virtualización del CPU y examinar planificación virtual del CPU en Xen. [Familiarity] • Esbozar la diferencia entre la clásica memoria virtual del SO y la virtualización de memoria. Explicar los múltiples niveles de mapeamiento de páginas en oposición a la virtualización de la memoria. Definir memoria <i>over-commitment</i> e ilustrar sobre Wmware <i>memory ballooning</i> como técnica de reclamo para sistemas virtualizados con memoria <i>over-committed</i>. [Familiarity]
Readings : [HDF11], [BVS13]	

Unit 5: Cloud Computing (12)	
Competences Expected:	
Topics	Learning Outcomes
<ul style="list-style-type: none"> • Cloud-based data storage <ul style="list-style-type: none"> – Shared access to weakly consistent data stores – Data synchronization – Data partitioning – Distributed file systems – Replication • Visión global sobre tecnologías de almacenamiento. • Conceptos fundamentales sobre almacenamiento en la cloud. • Amazon S3 y EBS. • Sistema de archivos distribuidos. • Sistema de bases de datos NoSQL. 	<ul style="list-style-type: none"> • Describir la organización general de datos y almacenamiento. [Familiarity] • Identificar los problemas de escalabilidad y administración de la big data. Discutir varias abstracciones en almacenamiento. [Familiarity] • Comparar y contrastar diferentes tipos de sistema de archivos. Comparar y contrastar el Sistema de Archivos Distribuido de Hadoop (HDFS) y el Sistema de Archivos Paralelo Virtual (PVFS). [Usage] • Comparar y contrastar diferentes tipos de bases de datos. Discutir las ventajas y desventajas sobre las bases de datos NoSQL. [Usage] • Discutir los conceptos de almacenamiento en la cloud. [Familiarity]
Readings : [HDF11], [BVS13]	

Unit 6: Modelos de Programación (12)	
Competences Expected:	
Topics	Learning Outcomes
<ul style="list-style-type: none"> • Visión global de los modelos de programación basados en cloud computing. • Modelo de Programación MapReduce. • Modelo de programación para aplicaciones basadas en Grafos. 	<ul style="list-style-type: none"> • Explicar los aspectos fundamentales de los modelos de programación paralela y distribuida. [Familiarity] • Diferencias entre los modelos de programación: MapReduce, Pregel, GraphLab y Giraph. [Usage] • Explicar los principales conceptos en el modelo de programación MapReduce. [Usage]
Readings : [HDF11], [BVS13], [Low+12], [Mal+10], [Bal+08]	

8. WORKPLAN

8.1 Methodology

Individual and team participation is encouraged to present their ideas, motivating them with additional points in the different stages of the course evaluation.

8.2 Theory Sessions

The theory sessions are held in master classes with activities including active learning and roleplay to allow students to internalize the concepts.

8.3 Practical Sessions

The practical sessions are held in class where a series of exercises and/or practical concepts are developed through problem solving, problem solving, specific exercises and/or in application contexts.

9. EVALUATION SYSTEM

***** EVALUATION MISSING *****

10. BASIC BIBLIOGRAPHY

[Bal+08] Shumeet Baluja et al. "Video Suggestion and Discovery for Youtube: Taking Random Walks Through the View Graph". In: *Proceedings of the 17th International Conference on World Wide Web*. WWW '08. Beijing, China: ACM, 2008, pp. 895–904. ISBN: 978-1-60558-085-2. DOI: 10.1145/1367497.1367618. URL: <http://doi.acm.org/10.1145>

- [BVS13] Rajkumar Buyya, Christian Vecchiola, and S. Thamarai Selvi. *Mastering Cloud Computing: Foundations and Applications Programming*. 1st. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2013. ISBN: 9780124095397, 9780124114548.
- [Cou+11] George Coulouris et al. *Distributed Systems: Concepts and Design*. 5th. USA: Addison-Wesley Publishing Company, 2011. ISBN: 0132143011, 9780132143011.
- [HDF11] Kai Hwang, Jack Dongarra, and Geoffrey C. Fox. *Distributed and Cloud Computing: From Parallel Processing to the Internet of Things*. 1st. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2011. ISBN: 0123858801, 9780123858801.
- [Low+12] Yucheng Low et al. “Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud”. In: *Proc. VLDB Endow.* 5.8 (Apr. 2012), pp. 716–727. ISSN: 2150-8097. DOI: 10.14778/2212351.2212354. URL: <http://dx.doi.org/10.14778/2212351.2212354>.
- [Mal+10] Grzegorz Malewicz et al. “Pregel: A System for Large-scale Graph Processing”. In: *Proc. ACM SIGMOD. SIGMOD ’10* (2010), pp. 135–146. DOI: 10.1145/1807167.1807184. URL: <http://doi.acm.org/10.1145/1807167.1807184>.