

SYLLABUS

1. Information regarding the programme

1.1 Higher education institution	Babeş-Bolyai University
1.2 Faculty	Faculty of Mathematics and Computer Science
1.3 Department	Department of Computer Science
1.4 Field of study	Computer Science
1.5 Study cycle	Master
1.6 Study programme / Qualification	High Performance Computing and Big Data Analytics

2. Information regarding the discipline

2.1 Name of the discipline		Operating Systems for Parallel and Distributed Architectures					
2.2 Course coordinator		Assoc. prof. Darius Bufnea					
2.3 Seminar coordinator		Assoc. prof. Darius Bufnea					
2.4. Year of study	1	2.5 Semester	1	2.6. Type of evaluation	E	2.7 Type of discipline	compulsory

3. Total estimated time (hours/semester of didactic activities)

3.1 Hours per week	4	Of which: 3.2 course	2	3.3 seminar/laboratory	1 sem + 1 pr
3.4 Total hours in the curriculum	56	Of which: 3.5 course	28	3.6 seminar/laboratory	28
Time allotment:					hours
Learning using manual, course support, bibliography, course notes					25
Additional documentation (in libraries, on electronic platforms, field documentation)					25
Preparation for seminars/labs, homework, papers, portfolios and essays					25
Tutorship					12
Evaluations					7
Other activities:					-
3.7 Total individual study hours	94				
3.8 Total hours per semester	150				
3.9 Number of ECTS credits	6				

4. Prerequisites (if necessary)

4.1. curriculum	<ul style="list-style-type: none"> • Operating Systems • Distributed Operating Systems • Computer Networks
4.2. competencies	<ul style="list-style-type: none"> • Average administration and programming skills

5. Conditions (if necessary)

5.1. for the course	<ul style="list-style-type: none"> • Video projector
5.2. for the seminar/lab activities	<ul style="list-style-type: none"> • Computers, Linux computers and Linux virtual machines for building a cluster, Network infrastructure

6. Specific competencies acquired

Professional competencies	<ul style="list-style-type: none"> • Capability of analysis and synthesis; • Understanding and working with basic concepts of data analysis and modelling; • Modelling and solving real-life problems; • Assimilation of mathematical concepts and formal models to understand the methods and components of high performance systems; • Capability of developing of high performance programs based on parallel and distributed programming; • Analysis, design, and implementation of data analysis systems; • Understanding and acquisition of methods of modelling, optimization, analysis of massive datasets, data visualization.
Transversal competencies	<ul style="list-style-type: none"> • Ethic and fair behaviour, commitment to professional deontology • Teamwork capabilities; able to fulfil different roles • Professional communication skills; concise and precise description, both oral and written, of professional results, negotiation abilities; • Entrepreneurial skills; working with economical knowledge; continuous learning

7. Objectives of the discipline (outcome of the acquired competencies)

7.1 General objective of the discipline	<ul style="list-style-type: none"> • Know the key concepts of parallel cluster architectures
7.2 Specific objective of the discipline	<p>At the end of the course, students will know how to</p> <ul style="list-style-type: none"> • build • deploy • configure • maintain • monitor • debug <p>a Linux parallel cluster</p>

8. Content

8.1 Course	Teaching methods	Remarks
1. Introduction to Operating systems for parallel architectures	Exposure, description, explanation, debate and dialogue, discussion of case	

	studies	
2. Parallel Cluster architecture: Cluster Head Nodes, Computer Nodes, Clustering Middleware	Exposure, description, explanation, case studies	
3-4. Parallel Cluster Paradigms: Single system image, Centralized system management, High processing capacity, Resource consolidation, Optimal use of resources, High-availability, Redundancy, Single points of failure, Failover protection and Disaster recovery, Horizontal and vertical scalability, Load-balancing, Elasticity, Run Jobs Anytime, Anywhere	Exposure, description, explanation, debate and dialogue, discussion of case studies	
5. Design and configuration. Network prerequisites for a parallel cluster: LAN, bandwidth, latency, interface, security aspects. Nodes automatic configuration and deployment	Exposure, description, explanation, case studies	
6. Virtualization of hardware, operating system, storage devices, computer network resources	Exposure, description, explanation, case studies	
7-8. Beowulf clusters deployment and administrations	Exposure, description, explanation, debate and dialogue, discussion of case studies	
9. Linux Cluster Distributions: Mosix, ClusterKnoppix. Automated operating systems and software provisioning for a Linux Cluster: Open Source Cluster Application Resources (OSCAR)	Exposure, description, explanation, case studies	
10. Cluster resources: distributed memory architecture and distributed shared memory, distributed file systems (examples: IBM General Parallel File System, Microsoft's Cluster Shared Volumes, Oracle Cluster File System)	Exposure, description, explanation, debate and dialogue, discussion of case studies	
11. Nodes and head node management, Cluster system management, Debugging and monitoring a parallel cluster, Node failure management	Exposure, description, explanation, case studies	
12. Data sharing and communication, Message passing and communication, Parallel processing libraries: Parallel Virtual Machine toolkit and the Message Passing Interface library	Exposure, description, explanation, case studies	
13. Software and development environment, Parallel application development and execution (Parallel Environment – PE), Job scheduling & management	Exposure, description, explanation, case studies	

14. Final review	Exposure, description, explanation, case studies	
Bibliography 1. Gregory Pfister: <i>In Search of Clusters</i> , Prentice Hall; 2 edition (December 22, 1997), ISBN-10: 0138997098, ISBN-13: 978-0138997090 2. George F. Coulouris, Jean Dollimore, Tim Kindberg: <i>Distributed Systems: Concepts and Design</i> , Addison-Wesley; 5 edition (May 7, 2011), ISBN-10: 0132143011, ISBN-13: 978-0132143011 3. Joseph D. Sloan: <i>High Performance Linux Clusters with OSCAR, Rocks, OpenMosix, and MPI</i> , O'Reilly Media (November 23, 2004), ISBN-10: 0596005709, ISBN-13: 978-0596005702 4. Daniel F. Savarese, Donald J. Becker, John Salmon, Thomas Sterling: <i>How to Build a Beowulf: A Guide to the Implementation and Application of PC Clusters</i> , The MIT Press (May 28, 1999), ISBN-10: 026269218X, ISBN-13: 978-0262692182 5. Gordon Bell, Thomas Sterling: <i>Beowulf Cluster Computing with Linux</i> , The MIT Press; 1 edition (October 1, 2001), ISBN-10: 0262692740, ISBN-13: 978-0262692748 6. Charles Bookman: <i>Linux Clustering: Building and Maintaining Linux Clusters</i> , Sams Publishing; 1 edition (June 29, 2002), ISBN-10: 1578702747, ISBN-13: 978-1578702749		
8.2 Seminar / laboratory	Teaching methods	Remarks
1. Project presentation	Conversation, debate, case studies	The Seminar/lab is organized as a total of 7 classes - 2 hours every other week
2. Cluster requirements	Conversation, debate, case studies	
3. Cluster building and deployment	Conversation, debate, case studies	
4. Cluster configuration	Conversation, debate, case studies	
5. Cluster maintenance	Conversation, debate, case studies	
6. Cluster debugging and monitoring	Conversation, debate, case studies	
7. Final evaluation of seminar/lab activities	Conversation, debate	
Bibliography Students, organized in teams of 4 or 5 members will have to build, deploy, configure, maintain, monitor and debug a Linux parallel cluster. The key concepts to accomplish these goals are presented during the course hours and are also available in the course' bibliography (see above).		

9. Corroborating the content of the discipline with the expectations of the epistemic community, professional associations, and representative employers within the field of the program

- Courses with similar content are taught for graduate students in major universities around the world,

including: Princeton, Berkeley, MIT.

- Course content is considered very important in the actual context of the increased need of computing power for computational science, interdisciplinary applications, and commercial applications as well, coupled with the high cost and low accessibility of traditional supercomputers.

10. Evaluation

Type of activity	10.1 Evaluation criteria	10.2 Evaluation methods	10.3 Share in the grade (%)
10.4 Course	- know the key concepts of parallel cluster architectures;	Written exam	30%
10.5 Seminar/lab activities	- know how to deploy, maintain, debug and monitor a parallel cluster	Presentation on a HPC related topic	30%
		Homework assignments	30%
		Default	10%
10.6 Minimum performance standards			
<ul style="list-style-type: none"> • At least grade 5 (from a scale of 1 to 10) at written exam and seminar/lab activities. 			

Date

Signature of course coordinator

Signature of seminar coordinator

Assoc. prof. Darius Bufnea

Assoc. prof. Darius Bufnea

Date of approval

Signature of the head of department

.....

.....