Identification	**Subject**	CMS 585 – Parallel Computer Architecture 4 KU/ 8 ECTS
	Department	Computer Science
	Program	Undergraduate
	Term	Fall, 2023
	Instructor	Shahnaz N. Shhabazova
	E-mail:	shahbazova@gmail.com
	Classroom/hours	41 Mehseti str. (Neftchilar campus), Tuesday 13:40-15:10 & Thursday 13:40-15:10
Prerequisites	Computer Architecture, Parallel Computing, Parallel Programs	
Language	English	
Compulsory/Elective	Required	
Required textbooks and course materials	***Core textbooks:*** Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University ***References:*** **2**. Parallel Computer Architecture: A Hardware/Software Approach (The Morgan Kaufmann Series in Computer Architecture and Design) 1st Edit 3. Computer Architecture: A Quantitative Approach (The Morgan Kaufmann Series in Computer Architecture and Design)	
Course outline	The presented training program presents the minimum requirements for students' knowledge and skills and defines the content and types of teaching methods on the subject "Architecture of parallel computing". It is intended for teachers and students who take part in the process of teaching the subject.	
Course objectives	The course "Parallel and Distributed Computing" offers the study of parallel and distributed architectures, algorithms for systems of linear equations and matrix inversions, iterative methods for solving nonlinear problems, shortest paths and dynamic programming, problems with network flow, fully asynchronous iterative algorithms, partially asynchronous iterative methods, as well as the organization of an asynchronous network of processors for distributed computing.	
Learning outcomes	Upon completion of this course, the students must be able to: Work on parallel and distributed computing covers several broad areas, such as the design of parallel machines, parallel programming languages, the development and analysis of parallel algorithms, as well as issues related to applications. The main focus is on algorithms, within this area we limit our attention primarily to numerical algorithms for solving systems of equations or optimization problems. We pay special attention to algorithms that allow a high degree of parallelization, such as Jacobi and Gauss-Seidel relaxation methods, and we comprehensively consider communication and synchronization issues. In particular, we discuss interprocessor interaction algorithms and conduct a comprehensive convergence analysis of asynchronous iterative methods. The development of parallel algorithms can	

<table>
<tr><td></td><td colspan="3">be carried out at several levels, and this to some extent explains the variety of literature on this issue.</td></tr>
<tr><td rowspan="4">Teaching methods</td><td colspan="2">Lecture</td><td>x</td></tr>
<tr><td colspan="2">Group discussion</td><td>x</td></tr>
<tr><td colspan="2">Experiential exercise</td><td>x</td></tr>
<tr><td colspan="2">Lab</td><td>x</td></tr>
<tr><td rowspan="6">Evaluation</td><td>Methods</td><td>Date/deadlines</td><td>Percentage (%)</td></tr>
<tr><td>Midterm Exam</td><td></td><td>30</td></tr>
<tr><td>Activity</td><td></td><td>10</td></tr>
<tr><td>Project</td><td>16/12/2023</td><td>20</td></tr>
<tr><td>Final Exam</td><td></td><td>40</td></tr>
<tr><td>Total</td><td></td><td>100</td></tr>
<tr><td>Policy</td><td colspan="3">

Preparation for class
The structure of this course makes your individual study and preparation outside the class extremely important. The lecture material will focus on the major points introduced in the text. Reading the assigned chapters and having some familiarity with them before class will greatly assist your understanding of the lecture. After the lecture, you should study your notes and work relevant problems and cases from the end of the chapter and sample exam questions.

Throughout the semester we will also have many review sessions. These review sessions will take place during the regularly scheduled class periods.
Withdrawal (pass/fail)
This course strictly follows grading policy of the School of Engineering and Applied Science. Thus, a student is normally expected to achieve a mark of at least 60% to pass. In case of failure, he/she will be required to repeat the course the following term or year.
Cheating/plagiarism
Cheating or other plagiarism during the Quizzes, Mid-term and Final Examinations will lead to paper cancellation. In this case, the student will automatically get zero (0), without any considerations.
Professional behavior guidelines
The students shall behave in the way to create favorable academic and professional environment during the class hours. Unauthorized discussions and unethical behavior are strictly prohibited.
Ethics
Students should not arrive in late to class.
All cell phones must be turned off and stowed away before entering class. Use of any electronic devices is not allowed in the classroom and violators will be punished accordingly.

</td></tr>
</table>

	• **Exams** All exams will be closed book. • **Project** This course is not devoted to the creation of software applications, but to the study of input-output control systems. The result of the project is the most important part, and the evaluation is given to projects that are executed on the computer. The number of people in each project group can be a maximum of 3 people. Participants have their own specific levels of task completion in the project. The presentation, which makes up a maximum of one-third of the entire project evaluation, should be presented in less than ten minutes, and answers to questions are also included. All team members are responsible for answering any questions about the project.

WK	Date/Day (tentative)	Topics	Textbook/Assignments
1.	28/09/2023 Lecture	Introduction. Why Parallel Architecture. Architectural Trends. Supercomputers. Communication Architecture. Data Parallel Processing. A Generic Parallel Architecture.	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 001
	28/09/2023 Laboratory	**Exercises**	
2.	05/10/2023 Lecture	Parallel Application Case Studies. Simulating Ocean Currents. Simulating the Evolution of Galaxies. Visualizing Complex Scenes using Ray Tracing. Data Mining for Associations.	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 2
3	12/10/2023 Lecture	The Parallelization Process. Steps in the Process. Parallelizing Computation versus Data. Parallelization of an Example Program. A Simple Example: The Equation Solver Kernel.	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 03
	12/10/2023 laboratory	**Exercises**	
4	19/10/2023 Lecture	Parallel Application Case Studies. Simulating the Evolution of Galaxies. Visualizing Complex Scenes using Ray Tracing. Mining Data for Associations. The Parallelization Process. Steps in the Process. Parallelizing Computation versus Data	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 04
5	26/10/2023 Lecture	Programming for Performance. Load Balance and Synchronization Wait Time. Reducing Inherent Communication. Data Access and Communication in a Multi-Memory System. A Multiprocessor as an Extended Memory Hierarchy. Artifactual Communication in the Extended Memory Hierarchy.	. Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 05
	26/10/2023 laboratory	**Exercises**	
6	02/11/2023 Lecture	Workload-Driven Evaluation. Scaling Workloads and Machine. Scaling Models. Scaling Workload Parameters. Evaluating a Real Machine. Choosing Workload. Evaluating a Fixed-Size Machin. Evaluating an Architectural Idea or Tradeoff. Multiprocessor Simulation.	. Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 06
7	09/11/2023	Shared Memory Multiprocessors. The Cache Coherence Problem. Cache Coherence	Parallel Computer Architecture. A Hardware / Software Approach

	Lecture	Through Bus Snoopin. Memory Consistency. Sequential Consistency. Sufficient Conditions for Preserving Sequential Consistency	David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 07
	09/11/2023 laboratory	**Midterm Exam** **Exercises**	
8	16/11/2023 Lecture	Design Space for Snooping Protocols. A 3-state (MSI) Write-back Invalidation Protocol. A 4-state (MESI) Write-Back Invalidation Protocol. Assessing Protocol Design Tradeoffs. Impact of Protocol Optimization. Tradeoffs in Cache Block Size.	. Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 08
9	23/11/2023 Lecture	Snoop-based Multiprocessor Design. Correctness Requirements. Base Design: Single-level Caches with an Atomic Bus. Cache controller and tags. Reporting snoop results. Base Organization. Non-atomic State Transition.	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 09
	23/11/2023 laboratory	**Exercises**	
10	30/11/2023 Lecture	Multi-level Cache Hierarchies. Maintaining inclusion. Propagating transactions for coherence in the hierarchy. Split-transaction Bus. Bus Design and Request-response Matching. Snoop Results and Conflicting Requests. Serialization and Sequential Consistency.	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 10
11	07/12/2023 Lecture	Case Studies: SGI Challenge and Sun Enterprise SMPs. SGI Powerpath-2 System Bus. SGI Processor and Memory Subsystems. SGI Challenge Memory System Performance. Sun Processor and Memory Subsystem.	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 11
	07/12/2023 laboratory	**Exercises**	
12	14/12/2023 Lecture	Scalable Multiprocessors. Bandwidth Scaling. Physical scaling. Scaling in a Generic Parallel Architecture. Realizing Programming Models. Primitive Network Transactions. Shared Address Space. Communication architecture design space.	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University Chapter 12
13	21/12/2023 Lecture	Physical DMA. User-level Access. ase Study: Thinking Machines CM-5. Dedicated Message Processing. Case Study: Intel Paragon. Shared Physical Address Space. Case study: Cray T3D.	Parallel Computer Architecture. A Hardware / Software Approach David Culler, University of California, Berkeley; Jaswinder Pal Singh Princeton University with Anoop Gupta Stanford University

			Chapter 13
	21/12/2023 laboratory	**Exercises**	
14	28/12/2023 Lecture	**Projects/Presentations**	
15	29/12/2023 Lecture	**Preparation for the final exam** **Exercises**	
	29/12/2023 laboratory	**Exercises**	