

High Performance Computing

Code: 102777
ECTS Credits: 6

Degree	Type	Year	Semester
2502441 Computer Engineering	OB	3	2
2502441 Computer Engineering	OT	4	2

Contact

Name: Anna Bàrbara Sikora
Email: anna.sikora@uab.cat

Use of Languages

Principal working language: catalan (cat)
Some groups entirely in English: No
Some groups entirely in Catalan: No
Some groups entirely in Spanish: No

Teachers

Sergio Villar Mesurado
Cristina Peralta Quesada
Anna Bàrbara Sikora

Prerequisites

Although there are no formally established prerequisites and the subject provides the students with a means to acquire the knowledge described in its content section, it is recommended a good knowledge of: programming, the functioning of a computer, different types of architectures, how an operating system works, how messages are sent through the network (Software Engineering, Computer Structure, Computer Architecture, Operating Systems, Networks, Advanced Architectures).

Objectives and Contextualisation

The objective of this subject is to know high performance and parallel computing systems, multiprocessor and multi-computing systems, paradigms of parallel programming, learn to develop applications using message passing or shared memory and analyze the performance of these applications.

The theoretical concepts about programming paradigms, message passing and shared memory are reinforced with lab sessions in which students learn to program using parallel programming languages.

All the components described in this subject must allow the student to understand the operation of high performance and parallel systems and, to some extent, be able to perform a simple design of a parallel application and evaluate its performance.

Competences

- Computer Engineering
- Acquire thinking habits.

- Communication.
- Conceive and develop centralised or distributed computer systems or architectures by integrating hardware, software and networks.
- Have the capacity to analyse and evaluate computer architectures, including parallel and distributed platforms, and develop and optimise software for the same.
- Have the capacity to analyse, evaluate, select and configure hardware platforms for the development and execution of computer applications and services.
- Have the capacity to design and implement system and of communications software.

Learning Outcomes

1. Analyse the available models, paradigms and parallel programming languages to determine which best fit the needs of an application.
2. Analyse the characteristics of applications that require high performance.
3. Apply the concepts of operating systems and networks to develop the necessary software components to manage a high performance computer system and the communications involved in these systems.
4. Communicate efficiently, orally or in writing, knowledge, results and skills, both in the professional environment and before non-expert audiences.
5. Configure hardware platforms and programming environments for the development and running of high performance applications.
6. Design, implement and integrate the software modules of the system and communications levels, considering the features of the same.
7. Determine the most suitable platforms for each type of application.
8. Develop a capacity for analysis, synthesis and prospection.
9. Develop parallel applications based on the existing paradigms.
10. Develop scientific thought .
11. Evaluate and predict the performance of different high performance computer platforms for running applications.
12. Evaluate the benefits of high performance architectures, as well as the functionality of their applications.
13. Evaluate the functionality and performance of developed parallel/distributed applications.
14. Identify the different types of high performance computing architectures, including parallel and distributed systems.
15. Identify the software levels that a high performance computer system should comprise.
16. Propose computer solutions based on high performance systems that integrate both the architecture of the system's hardware components and interconnection between the same, and the design of the necessary software.
17. Select the most suitable high performance computer system to satisfy the needs of a certain type of application or user.
18. Use efficiently ICT communication and transmission of ideas and results.

Content

Topic 1: Introduction to high performance systems.

Introduction to high performance systems, parallel systems, multiprocessors and multicomputers. Execution of parallel applications in high performance systems.

Topic 2: Classification of parallelism

SIMD (Single Instruction, Multiple Data), MIMD (Multiple Instruction, Multiple Data). Parallel application models.

Topic 3: Parallel Algorithms

Definition of parallel algorithms. Examples of parallel algorithms. Development of parallel algorithms.

Topic 4: Parallel programming

Paradigms of parallel programming. Applications based on message passing. Standard MPI (Message Passing Interface). Applications based on shared memory. Standard OpenMP (Open Multiprocessing) and CUDA (Compute Unified Device Architecture). Hybrid applications (MPI + OpenMP). Application development using MPI, OpenMP, CUDA and hybrid programming.

Topic 5: Architecture of Supercomputers

Particularities of the architecture of supercomputers. Interconnection networks. Physical infrastructure and costs.

Topic 6: Performance analysis

Analysis of performance of parallel systems. Evaluation of performance of parallel systems. Examples of performance evaluation tools.

Methodology

In the development of the subject, four types of teaching activities can be differentiated:

- Theoretical classes. Explanation on the board of the theoretical part of each theme of the program. The typical structure of a class of this type will be the following: in the first place an introduction will be made where the objectives of the class and the contents to be discussed will be briefly presented. Next, the contents studied will be uncovered, including narrative expositions, formal developments that provide theoretical foundations, and intercalating examples, which illustrate the application of the exposed contents. Finally, the professor will present the conclusions of the discussed contents. Throughout the course there will be continuous evaluations of the groups of topics.
- Practical classes. All the topics will be accompanied by a set of problems that the student must try to solve. In this sense, and as the student progresses in and deepens in their knowledge, these problems will be gradually more complex. The seminars will be the natural forum in which the development of practical work can be discussed in common, contributing the knowledge that the student is lacking to carry it forward. The mission of the practical classes is to bridge between theoretical classes and classes in the laboratory, which will promote the ability to analyze and synthesise, the development of critical reasoning, and that will train the student in the resolution of problems. Those that the professor considers of greater interest or in which the students find greater difficulty will be corrected on the board. Before the beginning of each set of problems the instructor may propose a list of exercises that students will have to solve.
- Classes in the laboratory. The practical part of the theoretical subjects will be completed with sessions in the laboratory, where the student will develop several programs and will have to try to solve a specific problem that the student will receive at the beginning of the semester. Some of these exercises will have to be delivered to the instructor at the indicated dates. The labs will be developed in groups of two students. The classes include 6 sessions in the laboratory, 2 hours long each, where the student will carry out the exercises.
- Deliverables. During the course, students (in groups) will do various deliverables. On the one hand, each group (of 5-6 people) will develop a research work (a scientific work) on a topic related to the subject. This exercise seeks to enable students to search, analyze and summarize available information (information search, scientific dissemination - books or articles) related to the concepts that are covered in the course. On the other hand, there will be a set of basic, short and practical exercises related to theory and practical classes of problems. Each group (2 people) will solve these exercises and make the delivery with the corresponding date. The aim of this exercise is for the student to be able to solve theoretically the concepts worked on in theory and then be able to apply them to the practical laboratory case.

This approach to work is aimed at promoting active learning and developing competencies of organizational and planning skills, oral and written communication, teamwork and critical reasoning. The quality of the exercises carried out, of its presentation and of its operation will be valued especially.

Transversal competences:

As mentioned before, the subject includes as a teaching activity a development of a special research project. It consists of research on a topic related to the subject and its presentation (e.g. a presentation on a scientific dissemination). The

work is aimed at developing students' capacity for analysis, synthesis and foresight and scientific thinking (competence T01, learning outcomes T01.02 and T01.03). Students will have guides that describe how to research the chosen topic. During the semester there will be meetings with the groups to discuss and evaluate the evolution of the work.

The management of the subject will be done through the Virtual Campus (<https://cv2008.uab.cat/>), which will be used to view the materials, manage the lab groups, make the corresponding deliveries, see the notes, communicate with the professors, etc.

Annotation: Within the schedule set by the centre or degree programme, 15 minutes of one class will be reserved for students to evaluate their lecturers and their courses or modules through questionnaires.

Activities

Title	Hours	ECTS	Learning Outcomes
Type: Directed			
Labs (Development of programming exercises in the lab)	9	0.36	1, 2, 11, 13, 12, 9, 8, 7, 6, 14, 16, 17
Problems (Solving exercises)	9	0.36	1, 2, 13, 4, 10, 8, 18, 16
Theory	20	0.8	1, 2, 3, 13, 5, 8, 14, 15, 17
Type: Autonomous			
Autonomous study	30	1.2	1, 2, 3, 7, 14, 15, 17
Labs preparation	40	1.6	1, 2, 11, 13, 12, 5, 9, 8, 7, 6, 14, 16, 17
Problem preparation	15	0.6	1, 2, 13, 4, 10, 8, 18, 16
Research	15	0.6	1, 3, 13, 5, 8, 7, 18, 14, 15

Assessment

The objective of the assessment process is to verify that the student has achieved the knowledge and skills defined in the objectives of the subject, as well as the associated competences.

Four types of activities will be assessed independently, and the weighted sum of them will give the final grade. These five activities are:

1. Theory (T)
2. Solution of the laboratory practices (PL)
3. Analysis of a research article (TR)
4. Completion of practical exercises (EXER).

The Theory (T) part will be assessed with two individual partial exams throughout the course. The final grade of Theory will come from the weighted sum of the two exams ($0.5 * \text{Control 1} + 0.5 * \text{Control 2}$). There will be a second chance to recover that part on the day at Juny we have assigned to. Parts that have not been passed in the partial theory exams may be recovered separately. The minimum grade to pass this part of Theory is $> = 5$. In the recovery only the failed exams can be recovered, the grade obtained cannot be raised (in the case of passing the Theory part). The maximum grade that can be obtained in recovery is 7.

The part of Laboratory exercises (PL) will be evaluated by group. There are three deliveries. The final grade will come from the weighted sum of the three deliveries ($X0 * \text{Delivery 1} + X1 * \text{Delivery 2} + X2 * \text{Delivery 3}$, $X0 + X1 + X2 = 1$). To pass the PL the minimum mark will have to be $> = 5$. There is only one opportunity to pass this part (this part can not be recovered).

The Research Work (TR) part will be evaluated in groups (5-6 people). Each group will have an assigned research topic on which to research, analyze, summarize, critique, and present. This work will be discussed with the group members at the indicated follow-up hours and subsequently a delivery will be made to the CV with the indicated date. The grade for each member of the group will be determined by the quality of the review and participation in the discussion. There is only one chance (this part cannot be recovered). Consequently, in the event that a student does not have a grade for the work, he / she will not be able to opt for the excellent mark in the final grade of the subject.

The practical exercises (EXER) will be done in groups of 2 people and will consist of working on very specific programming problems, short and related to those found in laboratory practices. We want each student to study the solution of a set of specific problems in isolation in the most general case of the practice. The value of these exercises is 20% of the final grade and given their nature and purpose are not recoverable.

The final grade of the subject will be the weighted sum of the grades of each of the four activities: 35% of Theory, 10% of Research Work, 20% Resolution of practical exercises and 35% of Resolution of laboratory exercises. The result will have to be ≥ 5 .

In case a student does not pass the subject due to not reaching the minimum score in any of the mandatory parts (Theory or Laboratory exercises), even though the weighted average is equal or superior to 5, the final grade of the subject will be 4.5.

In the event that the average does not reach 5, the official grade will be the average mark obtained numerically.

If the student delivers any activity, it is understood that he/she is participating in the subject and will be evaluated. If you do not deliver any activity, then it can be considered Non-evaluable.

Granting an honorific matriculation qualification is a decision of the faculty responsible for the subject. The regulations of the UAB indicate that MH can only be awarded to students who have obtained a final grade of 9.00 or more. It can be granted up to 5% of MH of the total number of students enrolled.

The dates of continuous evaluation and assignment delivery will be published on the virtual campus and may be subject to possible changes to adapt to possible incidents; the virtual campus will always inform about these changes since it is understood that the CV is the usual mechanism for exchanging information between professors and students.

For each assessment activity, a place, date and time of revision will be indicated in which the student will be able to review the activity with the professor. In this context, claims can be made about the activity grade, which will be evaluated by the professors responsible for the subject. If the student does not submit to this review, this activity will not be reviewed later.

SUMMARY

If $(T \geq 5) \wedge (PL \geq 5)$ then

$$NF = 0.35 * T + 0.1 * TR + 0.2 * EXER + 0.35 * PL$$

If $(NF \geq 5)$ then PASS

else FAIL

else FAIL

Repeating students: repeating students who have passed laboratory exercises in previous editions of the subject may request the validation of this part of the subject. The rest of the assessment activities must be carried out under the same conditions as the other students.

Note about plagiarism:

Without prejudice to other disciplinary measures deemed appropriate, and in accordance with the current academic regulations, irregularities committed by a student who may lead to a variation of the qualification in an assessable activity will be graded with zero (0). Assessment activities qualified in this way and by this procedure will not be recoverable. If it is necessary to pass any of these assessment activities to pass the subject, this subject will be suspended directly, without opportunity to recover it in the same course. These irregularities include, among others:

- the total or partial copy of a lab exercise, report, or any other evaluation activity;
- let another student to copy;
- present a group work not done entirely by the members of the group (applied to all members and not only to those who have not worked);
- present as own materials prepared by a third party, even if they are translations or adaptations, and generally works with non-original and exclusive elements of the student;
- have communication devices (such as mobile phones, smart watches, pens with camera, etc.) accessible during theoretical-practical assessment tests (individual exams);
- talk with classmates during the individual theoretical-practice tests (exams);
- copy or attempt to copy from other students during the theoretical-practical assessment tests (exams);
- use or attempt to use written material related to the subject during the theoretical-practical evaluation tests (exams), when these have not been explicitly allowed.

If you do not pass the subject due to the fact that some of the evaluation activities do not reach the minimum grade required, the numerical official grade will be the lowest value between 4.5 and the weighted average of the grades. With the exceptions that the "Non-Appraising" qualification will be awarded to students who do not participate in any of the assessment activities, and that the numerical official grade will be the lowest value between 3.0 and the average Weighted grades in case the student has committed irregularities in an evaluation act (and therefore the subject cannot be approved by compensation). In future editions of this subject, the student who has committed irregularities in an evaluation act will not be validated any of the assessment activities carried out.

In summary: copy, let copy or plagiarize (or attempt) in any of the assessment activities will lead to a FAIL, not compensable and without validations of parts of the subject in subsequent courses.

Assessment Activities

Title	Weighting	Hours	ECTS	Learning Outcomes
Individual partial exam 1	17,5%	2	0.08	1, 2, 8, 14
Individual partial exam 2	17,5%	2	0.08	1, 3, 9, 8, 7, 15, 16, 17
Practical exercises	20%	6	0.24	1, 11, 13, 9, 8, 7, 17
Programming exercise 1	8%	0.5	0.02	1, 3, 11, 12, 5, 9, 8, 6, 15, 16
Programming exercise 2	13,5%	0.5	0.02	1, 3, 11, 12, 5, 9, 8, 6, 15, 16
Programming exercise 3	13,5%	0.5	0.02	1, 3, 11, 13, 12, 5, 9, 8, 6, 15, 16
Research work	10%	0.5	0.02	1, 2, 12, 4, 10, 8, 7, 18, 14, 15, 17

Bibliography

- Parallel Programming Concepts and Practice. Bertil Schmidt, Jorge González-Domínguez, Christian Hundt, Moritz Schlarb. Morgan Kaufmann, 2018
- An Introduction to Parallel Programming, 2nd edition. Peter S. Pacheco. Morgan Kaufmann. 2018

- Programming Massively Parallel Processors, A Hands-on Approach. David B. Kirk, Wen-mei W. Hwu. Morgan Kaufmann. 3th Edition. 2018
- Computer Architecture. A quantitative approach. John L. Hennessy, David A. Patterson. Morgan Kaufmann. 6a edición. 2018
- Introduction to Parallel Computing. A. Grama et al. Addison Wesley, Second Edition, 2003.
- Parallel Program Development For Cluster Computing: Methodology, Tools and Integrated Environments. Edited by J. C. Cunha, P. Kacsuk, S. C. Winter. Nova Science Publishers, Inc., 2001.

Software

queue manager (SLURM)

tools for performance evaluation (perf, TAU, nvcc)

remote connection to the laboratory (e.g. MobaTex, Eclipse, Visual Studio, etc.)