

TCP P Curriculum: Adoption in the second year of the undergraduate Degree in Computer Science at the University of Murcia, Spain

Manuel E. Acacio, Joaquín Cervera, Javier Cuenca, Ricardo Fernández-Pascual, M. Carmen Garrido, Domingo Giménez¹, José Guillén, María-Eugenia Requena, Juan A. Sánchez

Courses involved: those in bold in Table I.

I. GENERAL LINES OF THE PROPOSAL

The proposal concerns the inclusion of parallelism topics at an initial level in the Degree of Computer Science (DCS) at the University of Murcia, Spain. Parallelism is taught during the second year of the degree, after the basic ideas of computing, mathematics and physics have been presented during the first year.

Up to now, some concepts of parallelism are being introduced in two subjects of the second year of the studies (Advanced Computer Architecture, **ACA**, and Concurrent and Distributed Programming, **CDP**), and our project proposes to extend the discussion of parallelism concepts to other subjects in the same year (Fundamentals of Operating Systems, **FOS**, and Algorithms and Data Structures, **ADS**). Two of the subjects (**FOS** and **ACA**) are taught in the first semester, and the other two (**CDP** and **ADS**) in the second semester. There are 10 different subjects in the curriculum of the second year of the **DCS**, so in our project parallelism will be tackled in 40% of the subjects and in the two semesters of the second year, which is the first level at which parallelism is treated.

Members of the three departments related with computing in the University of Murcia are collaborating in this proposal. **FOS** and **ACA** are taught by the Department of Computer Engineering (“Departamento de Ingeniería y Tecnología de Computadores”), **CDP** by the Department of Information and Communication Engineering (“Departamento de Ingeniería de la Información y las Comunicaciones”) and **ADS** by the Department of Languages and Systems (“Departamento de Lenguajes y Sistemas”). In addition, the Murcia Supercomputing Centre (**MSC**), belonging to the Science Park Foundation, is also collaborating in the project.

The applicants are responsible for some groups of these subjects: Joaquín Cervera and Domingo Giménez for **ADS**, M. Carmen Garrido and Juan A. Sánchez for **CDP**, Javier Cuenca for **FOS**, and Manuel Acacio and Ricardo Fernández-Pascual for **ACA**. Other lecturers or teaching assistants may join the experience once it begins. María-Eugenia Requena, head of the **MSC**, José Guillén, Project Manager, and some more members of the staff of the Mur-

cia Science Park Foundation will collaborate in the project depending on the necessities generated by the common experiences.

The main ideas of our project are:

- To introduce the concepts of parallelism at the first level in the **DCS**, in the **second year of the studies**, with a **multi-departmental** collaboration over **two semesters**. Joint experiences between pairs of subjects are planned.
- To extend the current treatment of parallelism in some subjects to more subjects, so as to **cover new topics** of the curriculum of the **TCP P**.
- To **introduce students to real-world situations** in which parallelism plays a major role. To do so, shared experiences with the **MSC** and practical experiences in the subjects involved in the project are envisaged.

The following sections treat the context of the project, the situation of parallelism in the Computer Science studies at the University of Murcia, and the proposed actions and outcomes of this project.

II. CONTEXT OF THE PROJECT

Computer Science studies in Spain have been recently reorganised. They have changed from a structure with three different degrees organised in three or five years to only one degree (“Grado en Ingeniería Informática”) of four years. In the previous organisation there were two lower degrees (three years) on Management (“Ingeniería Técnica en Informática de Gestión”, **ITIG**) and Systems (“Ingeniería Técnica en Informática de Sistemas”, **ITIS**) and one upper degree of five years (“Ingeniería Informática”, **II**). In this new degree, specialization is achieved toward the end with five different profiles (called ‘intensifications’) which comprise clusters of elective courses. Currently, the “Facultad de Informática” (Computer Science School) of the University of Murcia (**FI-UM**) offers five intensifications: Computer Engineering, Software Engineering, Computing, Information Technology and Information Systems. In the academic year 2010/2011 most of the Spanish universities have implemented the first or the first two years of the new degree, where we think that the basic concepts of parallelism should be included.

On the other hand, due to the advances in technology, parallel computing is becoming very popu-

¹Corresponding author, e-mail: domingo@um.es

lar. At present, not only clusters or supercomputers are parallel systems, but so are the standard computational systems. Laptops and desktops are multi-core and include GPUs that can be programmed to take advantage of parallelism. This caused the initiative of the IEEE Technical Committee on Parallel Processing to define the core topics of parallelism which should be included in undergraduate curricula, and now is a good moment to work on how parallel computing concepts should be included in the Computer Science degree in the Spanish university system.

As a consequence, a poster analysing the situation of the teaching of parallel computing in the Spanish university system and proposing some actions to include parallelism in core and specialised courses was presented in the EduPar11 workshop [1]. In that work, the situation in the universities where the authors work (Universities of Granada, La Laguna, Murcia and Polytechnic of Valencia) is analysed, and in this project we propose some actions for the inclusion of parallelism concepts at a first level in the DCS at the University of Murcia. Our project is guided by this analysis.

The University of Murcia is a medium-size university in Spanish terms, with approximately 30000 students. It is a generalist university, with scientific, health, humanist and social studies, but with only two engineering degrees (in Chemistry and Computer Science). Most of the engineering studies in the Región of Murcia are at the Polytechnic University of Cartagena. The project is for the Computer Science studies at the University of Murcia, and it is specific for the DCS in the FI-UM and for the course 2011-2012.

There are approximately 1200 students at the FI-UM. In the 2010-2011 course only the two first years of the DCS were implanted, and the DCS coexisted with the previous studies. In the second year of the studies, approximately 120 students were enrolled in the new degree and 140 remained in previous studies. Hence, second year students were enrolled in a variety of classes from different degrees. This made it difficult to take a common approach to the introduction of parallelism during the second year, which is the most appropriate year to do so in our opinion. A first experience was conducted during 2010-2011, but only with a small group of students. In the course 2011-2012, the DCS will be completely implemented, and the prevision is that most of the second year students will be enrolled in the DCS, organised in three groups of approximately 60 students, which in turn are divided in three subgroups of 20 students. Thus, the introduction of this experience in parallel computing teaching in the second year of the DCS during the course 2011-2012 will mean that most of the future professionals in Computer Science in the Región of Murcia will be initiated into parallel computing at an early stage, which we think is beneficial for their future careers and for the society where they will develop them.

The subjects in the second year of the DCS at the FI-UM are shown in Table I together with a brief description of their content. The complete structure of the DCS can be found in [2]. The subjects involved in this project are highlighted. They represent 40% of the credits in the second course, with two subjects in each semester, and comprise contents of systems and programming.

| First semester | |
|---|--|
| Algorithms and Data Structures I | |
| Algebraic Specifications, Hashing, Multiple and Dual Data Structures, Trees, Graphs, Search Algorithms | |
| Automata Theory and Computability | |
| Formal Languages, Grammars, Automata, Turing Machines, Computable Functions, Church-Turing Thesis, Limits of Algorithmic Computation | |
| Object-Oriented Programming | |
| Object-Oriented Design, Reuse and Maintenance, Classes, Objects, Inheritance, Polymorphism, Dynamic Linking, Unit Testing | |
| Fundamentals of Operating Systems, FOS | |
| Processes, Memory, Files, I/O, Security, Shell Scripts, Users Management, Filesystems, Backups, Monitoring | |
| Advanced Computer Architecture, ACA | |
| Performance Analysis, Pipelining, Control Dependencies, Static and Dynamic Scheduling of Instructions, Memory System Organisation and Performance | |
| Second semester | |
| Compilers | |
| Virtual Machines and Intermediate Languages, Lexical, Syntax and Semantic Analysis, Type Checking, Abstraction, Optimisation | |
| Databases | |
| Database Systems, Relational Database Systems: Model and Integrity Constraints, Relational Query Languages, Transactions | |
| Computer Networks | |
| OSI and TCP/IP Architectures, Data Link and Network Layers, LANs, WANs, Routing, Traffic Management, Socket Programming | |
| Algorithms and Data Structures II, ADS | |
| Analysis of Algorithms, Complexity, Greedy Algorithms, Backtracking, Branch & Bound, Game Trees, Divide and Conquer | |
| Concurrent and Distributed Programming, CDP | |
| Loosely and Strongly Coupled Systems Programming, Classic Programming Paradigms in Distributed Systems | |

TABLE I
SUBJECTS AND DESCRIPTORS IN THE SECOND YEAR OF DCS AT THE FI-UM.

III. PARALLELISM IN COMPUTER SCIENCE STUDIES AT THE UNIVERSITY OF MURCIA

The situation of parallelism in Computer Science studies at the University of Murcia (and other Spanish universities) and a proposal for inclusion of parallelism in the Spanish university system were analysed in [1]. Figure 1 depicts the current situation at the University of Murcia, and Figure 2 summarises the proposal. The figures represent the complete curriculum (compulsory courses in the first three years and specialisations in the fourth year), but our project is oriented to an initial introduction to parallelism during the second year, and so we summarise the situation for this year:

- No notions of parallelism are included in the first year of the studies, and we think it is better not to include it in the initial approach to computing in the first year.
- Some notions of parallelism are included in sys-

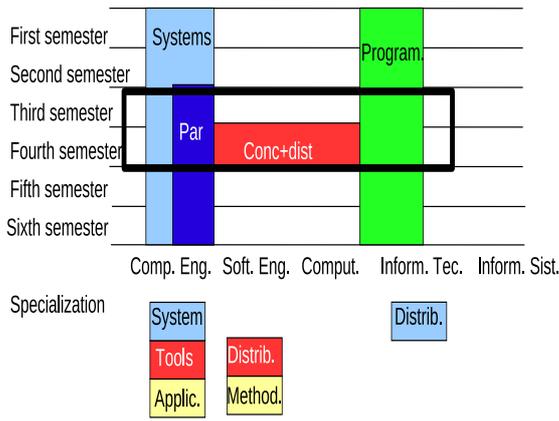


Fig. 1. Organisation of Parallel Computing teaching in the DCS at the FI-UM.

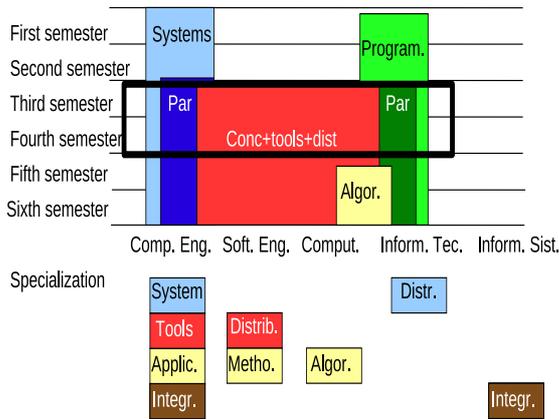


Fig. 2. Proposal of organisation of Parallel Computing in the Spanish DCS.

tems subjects (ACA), which we think is desirable. According to our project, notions of parallelism will be included in another systems subject (FOS).

- No notions of parallel programming or algorithms are considered in courses on sequential programming. In our project some notions will be included at least in one subject (ADS).
- There is a subject on the basic concepts of concurrency and distributed programming (CDP), but no applications or parallel systems are tackled in the laboratory sessions. The approximation to real applications and systems will be done through joint work with other subjects (FOS and ADS) and the MSC.

The topics of the TCPP curriculum covered in these subjects are shown in Table II. The second column contains how the topics appear in the TCPP curriculum. The letters in this table have the same meaning as in the TCPP curriculum: K stands for Know the term, C for Comprehend so as to paraphrase/illustrate, A represents Apply it in some way, and N Not in core. Currently, only two of the four subjects included in this project incorporate some concepts related to parallelism. Letters in black represent the current situation, and letters in red represent the situation we hope to have once our project is applied.

| TCPP | FOS | ACA | CDP | ADS |
|------------------------------|-----|-----|-----|-----|
| ARCHITECTURE | | | | |
| Architecture classes | K | | K | C |
| Superscalar (ILP) | K | C | | |
| SIMD/Vector | K | K | | |
| Pipelines | K | C | | |
| OoO execution | N | C | | |
| Multicore | C | | | |
| NUMA | N | K | | |
| Cache organization | C | | | |
| Atomicity | N | | C | |
| Impact memory hier. on soft. | N | K | A | K |
| Cycles per instruction (CPI) | C | C | | |
| Benchmarks | K | C | | |
| Spec marks | K | C | | |
| Peak performance | C | K | | |
| MIPS/FLOPS | C | C | | C |
| Sustained performance | C | K | | |
| PROGRAMMING | | | | |
| Shared memory | A | | C | A |
| Distributed memory | C | | C | K |
| Client server | C | | | |
| Task/thread spawning | A | | A | K |
| SPMD | C | | C | A |
| Shared memory notations | C | | K | A |
| Language extensions | K | | K | |
| Libraries | C | | A | |
| SPMD notations | C | | A | |
| MPI | C | | C | |
| Semantic tasks and threads | K | C | K | C |
| Synchronization | A | | A | |
| Critical regions | A | | A | A |
| Producer-consumer | A | | A | A |
| Monitors | K | | A | |
| Deadlocks | C | | K | |
| Memory models | N | | K | |
| Scheduling and computation | C | C | | |
| Decomposition strategies | C | K | | |
| Loop fusion | N | A | | |
| Scheduling and mapping | C | K | C | |
| Performance monitoring | K | A | | |
| Performance metrics | C | | C | A |
| Speed-up | C | | C | A |
| Efficiency | C | | C | A |
| Amdahl's law | K | | C | |
| ALGORITHMS | | | | |
| Asymptotics cost | C | | | C |
| Time | C | | | C |
| Space | C | | | C |
| Speedup | C | | | C |
| Notions from scheduling | C | | K | K |
| Divide and Conquer | C | | | A |
| Broadcast | C | | K | |
| Asynchrony | K | | K | |
| Synchronization | K | | A | |
| Sorting | C | | | A |
| Graph search | C | | | K |
| Specialized computations | C | | | K |
| CROSS CUTTING | | | | |
| Why and what is PDC | K | | C | |
| Concurrency | K | | C | |
| Non-determinism | K | | A | |
| Power | K | K | | |
| Locality | C | | C | |
| Security in Dist. systems | K | K | | |

TABLE II

TOPICS OF THE TCPP CURRICULUM IN PDC IN THE SECOND YEAR OF THE DCS AT THE FI-UM: IN BLACK THE PRESENT SITUATION, IN RED MODIFICATIONS AFTER APPLICATION OF THE PROJECT.

IV. ACTIONS AND OUTCOMES OF THE PROJECT

With this project parallelism topics will be introduced in two new subjects (FOS and ADS), and some new topics will be studied in the two subjects which already include parallelism in the present course 2010-2011. The new or modified topics appear in red in Table II. Furthermore, some activities are planned to enhance the coverage of parallelism topics, to coordinate how the topics are treated in the different subjects and to present students to real parallel applications and systems. The planned activities are represented in Figure 3. There are activities in individual subjects (activities 1, 5 and 6), involving two subjects (activities 3 and 4), and between two subjects and with the collaboration of the MSC (activities 2, 7 and 8). The activities are explained briefly below:

- **Act-1, FOS:** The concept of thread will be introduced, with some examples of use (overlapping

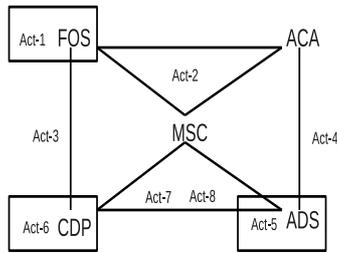


Fig. 3. Activities planned in the project.

of I/O and processing, and improvement of performance by using multithreading). Some of the questions of the implementation of threads will be analysed, comparing user and kernel implementations. Synchronisation problems will be outlined. In the laboratory sessions the use of the CPU when several threads are started will be analysed by using some system tools and profilers.

- **Act-2**, FOS+ACA+MSC: A visit to the Supercomputing Centre will be organised, together with a presentation of the staff of the Centre on practical aspects of a parallelism laboratory. The translation to a real situation of the topics studied in FOS and ACA will be analysed, and also some notions of cross cutting topics (why parallel computing, power, security, etc.) The activity will be scheduled at the end of the first semester.
- **Act-3**, FOS+CDP: In the second semester, when shared-memory and message-passing constructors are studied in CDP, the influence on the performance of the cores of the use of multiple threads and processes will be analysed. This will be done in collaboration with professors of FOS and applying the tools and concepts studied in the first semester.
- **Act-4**, ACA+ADS: In the second semester, once the analysis of sequential algorithms has been studied in ADS, the influence on the execution time of the memory hierarchy will be practically analysed in collaboration with professors of ACA and using the concepts studied in the first semester.
- **Act-5**, ADS: Seminars on algorithmic aspects of parallelism will be organised after some themes of ADS have been studied: parallel performance measures, after analysis of algorithms; and divide and conquer and dynamic programming parallel schemes, after the corresponding sequential schemes.
- **Act-6**, CDP: The basic shared-memory and message-passing constructors of OpenMP and MPI will be introduced so that they can be used for Act-7 and Act-8.
- **Act-7**, CDP+ADS+MSC: Shared-memory constructors studied in CDC will be used to implement some of the parallel algorithms introduced in ADS. Experiments will be carried out in the MSC with the Mooshak tool [3], which is used

in the practicals of ADS in the two semesters, and which has been adapted for parallel contests and installed in the MSC to use a cluster with four nodes, of eight cores each. The tool is currently being used in the first Spanish Parallel Programming Contest [4], and provides the speed-up achieved with OpenMP, MPI and hybrid programs.

- **Act-8**, CDP+ADS+MSC: Message-passing constructors studied in CDP will be used to implement some of the parallel algorithms in ADS. Experiments will be carried out in the system as for the previous activity.

The principal outcomes of the actions will be:

- The reinforcement of the introduction of parallel computing concepts to Computer Science students at an early level, by increasing the topics covered, the courses and lecturers involved in this experience, and by presenting students to practical aspects of parallel applications and systems.
- At the same time as students are introduced to parallel computing, new lecturers will join in, which should propitiate a better understanding of the importance of the subject, and consequently a better teaching and use of it, in parallel computing courses or in other subjects where parallelism can be useful.
- The collaboration in the second year of the studies of lecturers involved in higher level courses and in different specialities will allow us to prepare joint actions for successive courses at undergraduate level, and in the specialisations and in masters studies.
- Some material (documents, examples, practicals and tools), which could be used in future experiences at the University of Murcia or in other centres. Material will be generated for each one of the eight activities planned.

At the end of course 2011-2012, the action will be evaluated in two aspects and by two actors: to what extent each considered topic has been covered, and the usefulness of each activity will be analysed. The lecturers and students involved in the action will be questioned about their opinion (subjective) on the coverage of each topic and the success of each activity. For each activity, students will be evaluated as how much they have learned about each of the topics treated in the activity.

REFERENCES

- [1] F. Almeida, D. Giménez, J. M. Mantas, and A. M. Vidal, "On the integration of the Curriculum Initiative on Parallel and Distributed Computing in the Spanish university system," in *First NSF/TCPP Workshop on Parallel and Distributed Computing Education (EduPar-11)*, 2011.
- [2] BSc on Computer Engineering, University of Murcia, <http://www.um.es/informatica/index.php?pagina=bsc.on.computer.engineering.%282011-2012%29>.
- [3] Mooshak, System for managing programming contests on the Web, <http://mooshak.dcc.fc.up.pt/~zp/mooshak/>.
- [4] Spanish Parallel Programming Contests (in Spanish), <http://www.um.es/pcgum/eventos/cpp2011.html>.