

A learning experience toward the understanding of abstraction-level interactions in parallel applications



Alejandro Valero*, Rubén Gran-Tejero, Darío Suárez-Gracia, Emanuel A. Georgescu, Joaquín Ezpeleta, Pedro Álvarez, Adolfo Muñoz, Luis M. Ramos, Pablo Ibáñez

Department of Computer Science and Systems Engineering, Universidad de Zaragoza, Spain

ARTICLE INFO

Article history:

Received 30 October 2020

Received in revised form 22 February 2021

Accepted 18 May 2021

Available online 27 May 2021

Keywords:

Ray tracing

Task queue

Semaphore

Futex

Assembly instructions

ABSTRACT

In the curriculum of a Computer Engineering program, concepts like parallelism, concurrency, consistency, or atomicity are usually addressed in separate courses due to their thoroughness and extension. Isolating such concepts in courses helps students not only to focus on specific aspects, but also to experience the reality of working with modern computer systems, where those concepts are often detached in different abstraction levels. However, due to such an isolation, it exists a risk of inducing to the students an absence of interactions between these concepts, and, by extension, between the different abstraction levels of a system.

This paper proposes a learning experience showcasing the interactions between abstraction levels addressed in laboratory sessions of different courses. The driving example is a parallel ray tracer. In the different courses, students implement and assemble components of this application from the algorithmic level of the tracer to the assembly instructions required to guarantee atomicity. Each lab focuses on a single abstraction level, but shows students the interactions with the rest of the levels. Technical results and student learning outcomes through the analysis of surveys validate the proposed experience and confirm the students learning improvement with a more integrated view of the system.

© 2021 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The development of a Computer Engineering (CE) program must catch up with the fast evolution of the field. Since the end of the 2000s decade, technological limitations have led to an increase in the number of execution contexts running in parallel on a computer system. This requires the ability to not only implement algorithms that expose as much parallelism as possible, but also make an efficient use of hardware mechanisms to guarantee safe parallel execution. In this sense, following the recommendations of both the NSF/IEEE-TCPP Curriculum Committee [32] and the ACM/IEEE Joint Task Force on Computing Curricula [1], numerous approaches have made an effort to increase the presence or to reinforce Parallel and Distributed Computing (PDC) in CE programs. Recent work distributes PDC topics across different courses through the integration of modules into existing courses [7,20,38], introducing parallel programming in lower-level courses [6,18,42], the proposal of research-oriented teaching methodologies [16,22,30], or the creation of new courses [12,28,35].

A common approach to design and explain a computer system is to split the complexity of the whole system into self-contained levels. Since such levels relate to each other, each level provides a working interface to the remaining levels. These interfaces model a simplified abstraction of the underlying complexity and establish clear boundaries across the different parts of a system [21].

In most CE programs, each course typically resorts to abstractions in order to design and explain computer systems. Abstractions help to strengthen the learning process, since they make students focus on specific aspects. However, in our experience, students often lose the overall view of a computer system with such an approach. This may lead students to the conclusion that some courses are self-contained and do not relate to each other. Particularly, many of them forget the hardware implications underlying high-level abstractions, in terms of performance and power.

Previous work has proposed to teach PDC topics from the perspective of either high-level abstractions to ease both algorithm and software designs [13,17], or low-level abstractions such as assembly programming to understand what is required to support parallel execution [24]. Unlike these approaches, this paper reinforces PDC topics from the highest to the lowest level of abstraction that underlie complex parallel applications in a computer system [9]. More precisely, this work exposes to the students how the

* Corresponding author.

E-mail address: alvabre@unizar.es (A. Valero).

Instruction Set Architecture (ISA) and the operating system provide the required support to high-level synchronization operations, which in turn help strengthen the knowledge on how the essential concepts of parallelism, concurrency, consistency, and atomicity entangle among them and with the hardware [2,23,37].

To better understand the relations among the aforementioned concepts, this paper proposes to develop multiple components that at the end build a fully parallel ray tracing application. Ray tracing has been used in the past as a cross teaching experience to integrate two upper-level courses referring to high-level abstractions such as CUDA programming and advanced rendering concepts [25]. On the contrary, we present a learning experience involving multiple laboratory sessions of lower-level and upper-level courses of a CE program. The ray tracer serves as a motivating example that uses a concurrent queue to assign tasks to different execution threads. The queue is accessed in mutual exclusion to preserve data integrity. With this purpose, the access to the queue is managed according to each abstraction level, with *mutexes* implemented with library functions, system calls, or directly in assembly language. This way, the proposed learning experience covers four abstraction levels: Application, Library, Operating System, and ISA. Each abstraction level implicates a different course.

Each proposed lab is mainly tied to the interaction of two specific levels of abstraction, and purposely endowed with a context referring to the rest of the levels, contributing this way to integrate the different abstraction levels. In this work, we introduce the main guidelines, objectives, and results of the proposed experience, which allow to implement other experiences reinforcing inter-course learning.

Prior work has proved the suitability of a single-board computer for teaching parallel computing over mobile devices, student laptops, virtual machines, or remote multicore servers [26,27,45]. We build upon these studies by using a common hardware board in all the proposed labs, which contributes to consolidate an integrated view of the system. To this end, we analyze several boards and conclude that Raspberry Pi meets the vast majority of the hardware and software requirements of an inter-course learning experience.

The presented experience is the result of a project carried out in a CE program during the current and the past two academic years, in which assessment studies of the proposal have been already carried out thanks to a set of volunteer students. For the current 2020/2021 academic year, all the proposed labs are fully deployed and all the enrolled students in the involved courses are taking part in the proposed experience. This paper discusses experimental results for all the proposed labs, including both the technical details of the lab assignments and the students learning outcomes using pre/post surveys. These surveys expose that students effectively demand a deeper understanding of the interactions between the abstraction levels, and such demands are fulfilled after the completion of the labs.

The remainder of this paper is organized as follows. Section 2 introduces the context of the CE program and specific courses in which the proposed experience is established. Section 3 describes in depth the learning experience. Section 4 discusses the requirements to implement inter-course learning and the suitability of the selected boards. Section 5 shows the technical results. Section 6 presents a qualitative assessment of the applied learning method and students learning outcomes. Finally, Section 7 summarizes the paper.

2. Context

This section describes the organization of the CE program, including a brief description of the types of courses in each academic

year. In addition, the syllabus of the involved courses in the proposed learning experience is described in more detail.

2.1. CE program

The proposed experience is planned to be fully integrated in the CE program at the *Universidad de Zaragoza* (UNIZAR). This program consists of four academic years, 240 ECTS¹ credits in total.² The first two and a half years are common for all students. The core courses in this period mostly focus on the knowledge that any CE graduate should learn: algebra, calculus, discrete mathematics, programming theory, data structures and algorithms, computer architecture and organization, operating systems, physics and electronics, computer networks, databases, distributed systems, software engineering, artificial intelligence, and human-computer interaction. Afterward, students reinforce their knowledge in the major that most interests them within five available options: Computing, Computer Engineering, Information Systems, Information Technology, and Software Engineering. Each major consists of eight compulsory courses. In addition, students select two optional courses from any other major, as well as two core courses that are studied regardless of the chosen major. Finally, the students achieve the program by undertaking an undergraduate dissertation of 12 ECTS.

The CE program focuses on the application of theoretical knowledge in real-life problems, including the development of labs and projects. This approach is ideal to help students with the assimilation of the concepts studied in the different courses. Each course offers a well-thought out lab sessions tailored to reinforce the theoretical contents. At best, they are coordinated with other courses that belong to the same area of knowledge. As mentioned above, this can lead students to perceive a course, or a group of courses, as isolated islands, which makes it difficult for them to apply the knowledge acquired in each course in their professional career. In fact, these divisions are purely organizational, and all the courses have many interactions with each other. According to the Computer Engineering Curricula [2], students should learn the development of a *whole computer* in the lab experiments that include exposure to hardware and operating systems in the context of a relevant application, which is, in our case, the ray tracing algorithm.

2.2. Involved courses

The proposed experience implicates four different courses within the program to jointly face the problem. These courses are Computer Graphics, Distributed and Concurrent Systems Programming, Operating Systems, and Multiprocessors, which are related to the Application, Library, Operating System, and ISA abstraction levels, respectively.

Computer Graphics (CG) is a core course of the Computing major, and an optional course in other majors. CG focuses on mathematical models and algorithms that generate synthetic images (or videos) in which performance is a must. Students learn the underlying mathematical and physical concepts that define appearance and, as practical assignments, develop algorithms like ray tracing that output images from such concepts. Parallelization is key to the performance of such algorithms, and different parallelization strategies (static/dynamic, with different high level structures and partitions) must be explored.

Distributed and Concurrent Systems Programming (DCSP) is a core course that concentrates on the fundamentals of program-

¹ ECTS refers to European Credit Transfer and accumulation System: http://ec.europa.eu/education/resources-and-tools/european-credit-transfer-and-accumulation-system-ects_en.

² <https://estudios.unizar.es/estudio/ver?id=148>.

Table 1
Relations among the abstraction levels, courses, activities, academic years, semesters, and chronological order.

Abstraction level	Course	Activity	Academic year	Semester	Chronological order
Application Library	CG	Ray tracing	4 th	Fall	4 th
	DCSP	Concurrent task queue	2 nd	Fall	1 st
Operating System	OS	Futex system calls	2 nd	Fall	2 nd
ISA	MP	Futexes with assembly code	3 rd	Spring	3 rd

ming such classes of systems. In the case of concurrency, the lectures focus on the explanation of the problems that arise when a set of processes have to share data and resources, and the way such a problem has been solved, from the main mutual exclusion algorithms based on shared variables to higher level structures such as semaphores and monitors. In the case of distributed systems, students learn how to coordinate processes by means of synchronous and asynchronous message passing as well as by means of the use of a shared tuple space. Besides studying the concepts from a conceptual point of view, students work in a set of laboratory sessions and a final team project in which they have to develop some programs where the studied concepts are a crucial part.

Operating Systems (OS) is a core course that presents in a comprehensive way the structure and functions of an operating system. The operating system is presented as a resource manager and as a service provider at the system call and command interpreter levels. At each level, the student acquires concepts and skills related to the management and the use of the main system resources such as the processor, memory, and input/output devices. In relation to the topic presented in this paper, the course presents the synchronization primitives offered in the pthread library and studies the keys to their implementation. First, it analyzes the implementation of spinning primitives with the support offered by the processor in the form of atomic memory access instructions and then it motivates the need of the operating system support to implement sleeping primitives.

Multiprocessors (MP) is a core course of the Computer Engineering major, and an optional course in other majors. MP focuses on the mechanisms that support the parallel execution of tasks in a computer system from the point of view of the architecture and the organization of a computer. More precisely, this course focuses on parallel processors with shared memory, which are basic elements of current complex digital systems. The covered topics include performance analysis, performance modeling, on-chip networks, atomicity, consistency, and coherence in the memory hierarchy of a parallel processor. In the laboratory sessions of the course, OpenMP is presented and used as a tool for parallel programming of computers, as well as tools for performance measurement.

3. Proposed learning experience

This section presents the proposed experience that helps students to accomplish an integrated view of a computer system. The lab materials and resources for each abstraction level consist of a description of the work to be done, code snippets, and a series of milestones, where each one builds on top of the previous one. The experience involves a total of eight hours, since each lab session comprises two hours in the course associated with the level. Interested readers may refer to the following repository with the source code of every lab: <https://github.com/universidad-zaragoza/learning-experience-ray-tracing>.

3.1. Overview

The proposed learning experience allows students to consolidate the concepts of parallelism, concurrency, consistency, and atomicity exploitable in current multicore computers. We focus on ray tracing, an appealing application which can be efficiently parallelized by learning and using the above concepts. Table 1 shows the involved four levels of abstraction and the associated courses. The table also shows, for each course, the academic year, semester, and chronological order in which the activities take place.

According to the chronological order, students start the experience in the second academic year. The first lab, which belongs to the DCSP core course, focuses on the library level. This lab deals with the implementation and management of a task queue with concurrent access by multiple threads. Synchronization aspects must be considered in order to avoid race conditions. To do so, students use a semaphore library for handling such synchronization questions.

The subsequent lab takes place shortly, during the same academic year and semester, and focuses on the Operating System level. In this core lab, a mutex is implemented with a futex (fast userspace mutex) mechanism through atomic primitives and operating system calls that are only invoked when the mutex is contested [14]. This mutex is then used to implement a new semaphore library, which replaces the one used in the previous lab.

The following academic year covers the third lab, that is, the Assembly level, which is developed in the optional MP course. In this lab, assembly instructions are used to implement the mutex/futex, which have the potential to achieve a greater efficiency in energy consumption and performance compared to library functions and system calls.

Finally, in the fourth year, the students focus on the Application level by implementing a ray tracer in a lab of the CG optional course. In this activity, the rendering of an image is parallelized by dividing the image into regions. These regions are assigned to different threads by using the concurrent task queue. At this moment, the students fully evaluate and state the differences of protecting concurrent accesses to the task queue by using library functions, system calls, or assembly instructions.

Note that the development of the presented experience is subject to certain risks; e.g., students transferring from one institution to another, or students failing a course or simply not choosing the involved optional courses would not complete the full experience. To mitigate such risks, all the labs include two parts. The first one, which is self-contained, includes the material for the actual lab, and the second part links the lab with the others. Therefore, if a student does not complete a preceding or following lab assignment, the faculty can provide a solution, so that students can accomplish the second part of the lab and establish the links between the abstraction levels.

3.2. Abstraction levels

The application under study is presented in the next sections following the chronological order that students will experience.

3.2.1. Concurrent task queue

The aim of this lab is the implementation of one of the most common concurrent data structures: a queue. Queues, whose sequential approach has already been studied in a previous course of Data Structures and Algorithms, are a very suitable mechanism for the collaborative work of a set of processes. In the considered case, producers and consumers can, in a natural way, use one or more queues to share information and to synchronize [44]. As in any shared data structure, in order to preserve data integrity, the concurrent access to the shared data requires the use of some synchronization mechanisms.

In this context, the main objectives of this lab work are as follows: i) to implement a concurrent bounded queue for generic data types, ii) to get familiarity with semaphores as a mean for solving synchronization problems, and iii) to use a general and powerful approach to solve general synchronization problems using semaphores.

Controlling the concurrent access to a queue requires to consider not only mutual exclusion access to some queue components, but also condition synchronization (no first element exists in an empty queue, or no new element can be inserted when the queue is full). According to the focus proposed for the DCSP course, as a first assignment, students have to design the concurrent access to the queue using the coarse-grain atomic statement $\langle \text{await } B \ S \rangle$, where B is a boolean guard, usually concerning shared data, and S is a block of sequential statements. The semantics of the statement ensures that S starts its execution being B true, and the whole statement is atomically executed. The high-level point of view of such a statement makes easier the task of designing correct concurrent programs, which is one of the aims of the course. This will be done in the generic class `ConcurrentBoundedQueue` sketched in Listing 1, which includes a `BoundedQueue` as one of its attributes. Students have to code the complete data structure, including both enqueue and dequeue operations using semaphores for synchronization.

```

template <class T>
class ConcurrentBoundedQueue {
public:
    void enqueue(const T d);
    // <await this->bq->length()<N
    //   this->bq->enqueue(d)
    // >

    void dequeue();
    // <await this->bq->length()>0
    //   this->bq->dequeue()
    // >
    ...
private:
    int N; //size of the bounded queue
    BoundedQueue<T> *bq; //data storage
    ...
};
    
```

Listing 1: `ConcurrentBoundedQueue` generic class.

In previous lectures, students have seen the *pass the baton* technique [3] to implement $\langle \text{await } \dots \rangle$ statements using (binary) semaphores. This technique works as follows. Let us consider all the $\langle \text{await } B_i \ S_i \rangle$ and $\langle S_j \rangle$ sections that must be synchronized together. A mutex semaphore ensures exclusive access to such code areas. For each different B_i , a (binary) semaphore b_i with an initial count equal to 0 is added to block a process in the $\langle \text{await } B_i \ S_i \rangle$

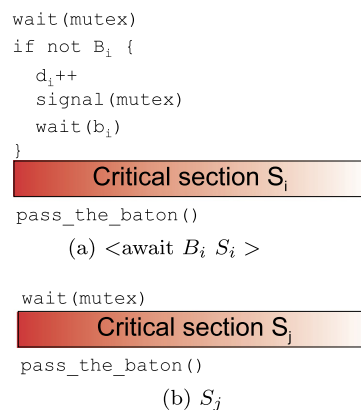


Fig. 1. Implementation of critical sections using (binary) semaphores.

statement when B_i is false. In addition, a counter d_i with an initial value set to 0 is included to store the number of processes blocked at B_i . Fig. 1 illustrates an implementation of both $\langle \text{await } B_i \ S_i \rangle$ and $\langle S_j \rangle$ sections. The *pass the baton* implementation is sketched in Listing 2.

```

void pass_the_baton() {
    switch {
        ...
        Bi and di>0:
            di --
            signal(bi)
        ...
        otherwise:
            signal(mutex)
    }
}
    
```

Listing 2: `pass_the_baton` function.

In the case of the concurrent bounded queue, there are only two conditions: non-empty and non-full. Therefore, three semaphores and two counters will be used. The private part of the `ConcurrentBoundedQueue` class can be extended as shown in Listing 3.

```

template <class T>
class ConcurrentBoundedQueue {
public:
    ...
private:
    int N; //size of the bounded queue
    BoundedQueue<T> *bq; //data storage
    Semaphore *mutex; //initial count will be 1
    Semaphore *b_not_full; //to block until queue is not full
    Semaphore *b_not_empty; //to block until queue is not empty
    int d_not_full; //# of b_not_full-blocked processes
    int d_not_empty; //# of b_not_empty-blocked processes
    void pass_the_baton(); //to pass the baton
};
    
```

Listing 3: `ConcurrentBoundedQueue` class including a synchronization mechanism.

Since we want students to finally implement general semaphores, the presented technique will be developed using general semaphores instead of binary semaphores. Anyway, notice that the way the technique is implemented, every semaphore count is always 1 or 0, having a behavior equivalent to binary semaphores. At this level, the (binary) semaphores are the lowest abstraction construct to manage synchronization, considered as an abstract data type. The students know two possible semantics, equivalent from

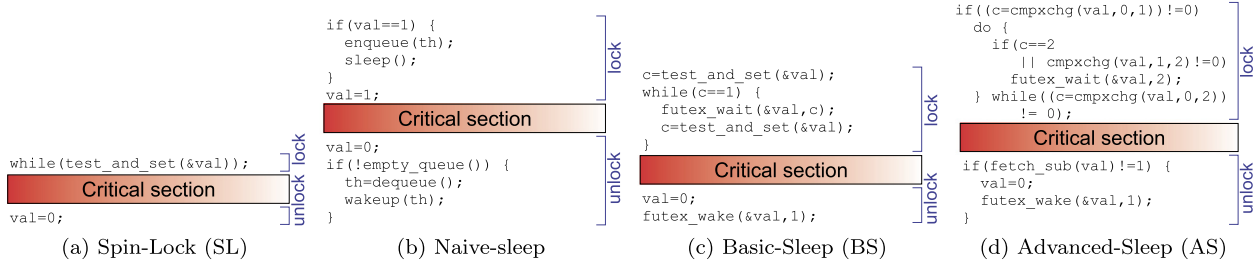


Fig. 2. Lock and unlock procedures of spin-lock and sleep mutexes. The sleep implementations include operating system calls to change the thread status.

a safeness point of view, but with possible different liveness properties. However, they still do not know how a semaphore can be implemented and how operates internally, whether it makes a process spin until the access is granted or the process goes to sleep controlled by the operating system. These issues are outlined in the lab, and students will find out the answers by implementing the semaphore abstract data type in the two following labs.

As a final and optional lab assignment, students are asked to implement a second approach of the concurrent queue. In the first one, each operation on the queue is executed in mutual exclusion. In the second one, students have to adapt the readers-writers approach so as to allow multiple access to *reading* operations (operations with no side effects on the queue) while preserving mutual exclusion access for *writing* operations, giving priority to writers in case of conflict. During the lectures, students have already designed a solution based on the `<await ...>` instruction, whereas this assignment deals with the implementation.

After completing this lab, students will have reinforced their knowledge about the main concepts related to semaphore-based synchronization. In addition, the proposed assignments also deal with the use of design techniques focusing on the synthesis of correct concurrent programs.

3.2.2. Task queue protection with futex system calls

This lab is intended to present the mechanisms required by the operating system to provide synchronization in concurrent algorithms. The main objectives of this lab are: i) show the operating system as a service provider for the user through system calls, ii) learn an efficient use of the futex system calls and the primitives of atomic instructions provided by the operating system and the C standard library, iii) understand the necessary mechanisms to provide execution in mutual exclusion with futexes and atomic instructions, and iv) show and use self-implemented lock and unlock primitives of a mutex abstraction to manage the access to the concurrent task queue implemented in the previous activity.

The lab material firstly describes the C11 atomic instructions from `stdatomic.h` and solicits the students to implement a mutex with spin-lock based on atomic instructions. Next, the *sleep* approach of a mutex is motivated, introducing the mandatory intervention of the operating system to change the thread status, and providing a naive approach of the sleep mutex using hypothetical `sleep` and `wakeup` system calls as well as management operations on a system queue. The limitations of this approach are used to motivate the futex system calls. Then, the syntax and use of the parameters of the `futex_wait` and `futex_wake` system calls are described. By using these calls, the students are guided to implement an intuitive and straightforward version of the sleep mutex referred to as basic implementation. Finally, the pseudo-code algorithm of a more efficient mutex is offered as a guideline to code an advanced implementation. This approach is based on the mutex implementation proposed by U. Drepper [10], which is integrated into the Linux kernel [11].

Fig. 2(a) shows the lock and unlock procedures of a Spin-Lock (SL) mutex protecting a critical section. The value of the userspace

`val` variable represents the two states of the mutex: not taken (`val=0`) and taken (`val=1`). The `test_and_set` atomic instruction changes the mutex state.³ More precisely, this instruction sets `val` to 1 and loads its previous value into `c` without the overhead of a system call. Then, a thread enters into the critical section if the lock is uncontested (`c=0`). Otherwise, the thread keeps spinning in the lock. In the unlock procedure, the thread simply sets `val` to 0 to release the mutex. Since the SL mutex leaves all the waiter threads in the lock awake, it may suffer system performance losses when the mutex is contested.

Fig. 2(b) illustrates the naive-sleep approach of a mutex. These procedures are similar to other versions offered in textbooks of operating system concepts such as [36,39], and [5]. This code is only correct if both procedures are executed atomically. However, assuming a non-atomic execution presents several problems that are listed in the lab material and should be understood by the students, specifically: i) the reading and writing operations of `val` are not atomically performed, which can lead to multiple threads reading the lock as not taken, ii) the reading of the lock and the insertion of the thread in the queue are neither atomic, which can lead to an indefinitely suspended thread if the lock is freed between the reading and insertion operations, and iii) after waking up from the `sleep` call, a thread has no guarantee of obtaining the lock in mutual exclusion since another thread can enter into the critical section before the former takes the lock.

Fig. 2(c) shows the Basic-Sleep (BS) implementation addressing all the incorrect behaviors stated above. In the lock function, the atomic operation changes the state of the mutex. If the lock is uncontested, the kernel is not invoked and the thread enters into the critical section. Otherwise, the `futex_wait` system call is invoked. It suspends the calling thread in a system queue if the lock is still taken (`val=1`), or it returns immediately if the lock has been released in the meantime (`val=0`). In the first case, the thread remains suspended until another thread wakes it up. Notice too that every time `futex_wait` returns, the thread tries to acquire the lock again.

The unlock procedure sets `val` to 0 and calls `futex_wake`. This system call wakes up a number of threads stated in the second argument (1 in the example as only a single thread is allowed to enter into the critical section) from those suspended in the system queue. Notice that such a call is invoked regardless of the lock is uncontested or not, which may impact on the system performance.

The Advanced-Sleep (AS) implementation shown in Fig. 2(d) addresses the performance problem of the basic approach. In this case, there are three mutex states: not taken (`val=0`), taken and no waiter threads (`val=1`), and taken and at least one waiter thread (`val=2`). In the lock procedure, `test_and_set` is no longer useful since `val` takes three values. Instead, the atomic `cmpxchg` primitive is used, in which a 1 (*desired* third argument) is loaded into

³ For the sake of brevity, we have shortened the original `stdatomic.h` function names; e.g., `test_and_set` corresponds to `atomic_flag_test_and_set` and the assignment operator for `val` refers to `atomic_store`.

val on a successful comparison between val and 0 (expected second argument). Regardless of the result of the comparison, the original value of val is loaded into c. If c==0, the calling thread updates the state of the mutex as taken and no waiters, and then enters into the critical section. Otherwise, the thread is suspended in the system queue by calling futex_wait. Previously, the second cmpxchg sets val to 2 if necessary, updating the state of the mutex as taken and at least one waiter. Note that, if the lock is freed between the first and second cmpxchg, the latter returns 0 and the thread is not suspended. The third cmpxchg ensures that a thread takes the mutex only if a 0 is returned. In such a case, val is set to 2 because there is no certainty of the number of waiters.

The unlock method subtracts 1 to val with the atomic fetch_sub, which returns the previous value of the argument. The futex_wake call is invoked just in the case of a suspended thread in the lock, avoiding such costly system calls when there are no waiter threads. The reader is referred to [10] for further details about the AS mutex implementation.

Once the different mutexes have been coded and understood, the students use them to support a complex abstraction, that is, the concurrent task queue implemented in the previous lab. Listing 4 shows an implementation alternative of the Semaphore class introduced in Listing 1, referred to as Library. This approach uses standard library mutexes (std::unique_lock <mutex>) to ensure mutual exclusion in Semaphore class methods. On the other hand, Listing 5 shows a different implementation of the same class, referred to as Thread-suspension, in which the lock and unlock methods are replaced by the procedures of each mutex version (see Fig. 2).

```
class Semaphore {
private:
    std::mutex mtx;
    std::condition_variable_any cv;
    int count = 0;
public:
    ...

    void Semaphore::signal() {
        std::unique_lock<mutex> lck(mtx);
        count = 1;
        cv.notify_all();
    }

    void Semaphore::wait() {
        std::unique_lock<mutex> lck(mtx);
        while (count==0)
            cv.wait(lck);
        count = 0;
    }
};
```

Listing 4: Library implementation alternative of the Semaphore class.

For both implementation alternatives, the wait method consists of a loop over a count variable, but, if count equals to 0, the current thread suspends its execution. On the other hand, the signal method wakes up all suspended threads after freeing the semaphore. In order to implement such a functionality for Thread-suspension, the suspend and wakeup methods are coded by the students using futex system calls. On the contrary, the Library approach relies on condition variables to suspend/wake up threads. Using both alternatives, the students assess the suitability of not only their coded spin-lock and sleep mutexes, but also a library mutex by measuring the execution time and the chip temperature under different contention scenarios (see Section 5.2).

Overall, the students will be able to use futex system calls and atomic instructions to implement spin-lock and sleep versions of

```
class Semaphore {
private:
    std::mutex mtx;
    int count = 0;
public:
    ...

    //suspends thread iff ve!=count
    void Semaphore::suspend(int ve) {
        syscall(__NR_futex, &(count), FUTEX_WAIT, ve, NULL, 0, 0);
    }

    //wake up all suspended threads
    void Semaphore::wakeup() {
        syscall(__NR_futex, &(count), FUTEX_WAKE, INT_MAX, NULL,
0, 0);
    }

    void Semaphore::signal() {
        mtx.lock();
        count = 1;
        wakeup();
        mtx.unlock();
    }

    void Semaphore::wait() {
        mtx.lock();
        while (count==0) {
            int vr = count;
            mtx.unlock();
            suspend(vr);
            mtx.lock();
        }
        count = 0;
        mtx.unlock();
    }
};
```

Listing 5: Thread-suspension implementation alternative of the Semaphore class. Mutex lock and unlock procedures refer to the different approaches from Fig. 2.

a basic synchronization abstraction such as a mutex, incorporate such approaches to protect the concurrent task queue, and experimentally state the performance differences among them.

3.2.3. Futexes with assembly code

The main purpose of this lab is to help students understand the support provided by the ISA level to implement fast and reliable mutual exclusion, in terms of consistency and atomicity. The ARM processors include load-link/store-conditional instructions and memory barriers, providing the foundation for higher level structures such as mutexes and futexes. In addition, these instructions do not require any privilege level for being executed, so programmers can directly exploit them to improve efficiency and reduce the overhead of systems calls.

By the end of this lab, students will have accomplished the following goals: i) understand how atomic instructions operate at the ISA level for the ARMv8 processors, ii) know why data memory barriers are often required when writing atomic instructions, and iii) learn the performance and energy implications of the different mutex implementations.

The assignments of this lab are designed to help students to engage with complex code enhancing their low-level programming skills, especially concerning performance and energy efficiency. In addition, they show how important is for an ISA to provide support for complex high-level constructors such as the mutexes used by operating systems, libraries, and applications. Finally, students gain knowledge on the relation between the C/C++11 memory model and the corresponding consistency models at the ISA level.

The lab material of this session is organized in two parts. In the first part, the students are asked to generate a race condition with the writing of a multi-threaded program that reduces an array by

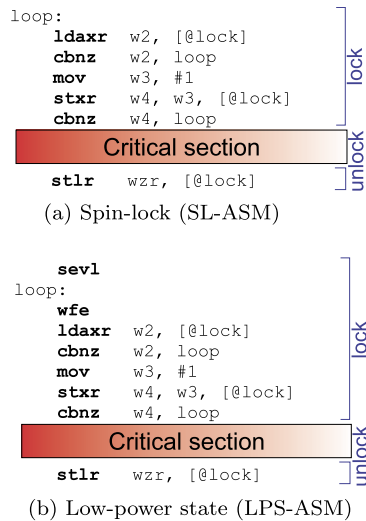


Fig. 3. Lock and unlock procedures with ARMv8 assembly code.

adding all the elements without synchronization primitives. Then, the students code a *fetch and add* primitive with ARMv8's *load-link* (*ldaxr*) and *store-conditional* (*stlrr*) instructions [4]. The implemented *fetch and add* is included in the previous program to verify that the code is now free of race conditions.

The second part comprises two assignments. The first one proposes a basic implementation of lock and unlock mutex functions based on *ldaxr/stlrr* instructions as plotted in Fig. 3(a). Threads in the lock function spin until they acquire the lock. This mutex approach is referred to as SL-ASM. The spinning can occur at the two *cbnz* instructions. Either if the lock is already taken (first *cbnz*) or the *stxr* instruction fails the attempt to take the lock (second *cbnz*), the branch instructions return the flow to the beginning of the loop. Notice too that, likewise the SL and BS implementations from the OS level, just two mutex states, taken and not taken, are considered in the assembly level.

The second assignment proposes an advanced implementation of the lock function by replacing the power-hungry spin-lock with a *wfe* instruction. This instruction puts the core into a low-power state without returning the control to the operating system. Fig. 3(b) shows such an energy-efficient implementation, referred to as LPS-ASM, also with the two mutex states taken and not taken. The student will learn how the operating system considers that the program is running, while it is actually waiting for the lock to be released, and how the thread can regain the lock without a system call. In particular, the *stlrr* instruction, located in the unlock function, performs a store with a release barrier and wakes up any core that could be in a low-power state after executing a *wfe* instruction. To guarantee progress, the cores also leave the low-power state after an interruption occurs (e.g., a context switch).

With both SL-ASM and LPS-ASM implementations, students will carry out a quick comparison between them in terms of performance and energy consumption. Since Raspberry does not provide energy hardware counters, as an indirect measurement, we periodically measure the temperature provided by the chip itself, which is mapped by the OS in the filesystem.⁴

3.2.4. Parallel ray tracing

The CG course proposes a practical assignment involving the implementation of a ray tracing algorithm [43], which is parallelized by assigning different tasks (partitions or regions of the




Fig. 4. Diagrams of a 2D image split into render tasks with different kinds of image regions and sizes.

expected synthesized image) to different threads. The main objectives of this lab are: i) find and understand the computational bottlenecks of the algorithm, ii) devise parallelization strategies that affect performance without any accuracy loss, and iii) test, explore, and analyze the impact (and overhead) of the combination of different parallelization strategies, including partitioning structures and thread assignment methodologies, on performance.

The contents of this lab include a description of the ray tracer, and an introduction on how to parallelize it. In particular, this assignment makes use of the minimalist C++ *smallpt* ray tracer by K. Beason.⁵ This algorithm generates a 2D image from a 3D representation of a virtual scene, including geometry and optical properties of the objects and physical characterizations of sensors (cameras) and light sources. In practice, the algorithm simulates light transport paths across the virtual scene in order to obtain the final color that reaches each of the pixels of the image. Paths are generated from the camera and traverse each pixel independently. Since the computation associated to each pixel is independent, the algorithm is highly parallelizable. Moreover, such a parallelization is worthwhile because the algorithm is computational intensive and takes quite a long time to converge (about 1 or 2 hours for a good quality result for a simple virtual scene, and even days in the case of more complex scenes).

A common ray tracing parallelization strategy is to subdivide the image into different regions, converting the computation of each of the regions into a render task to be assigned to an execution thread. The students are required to explore different parallelization strategies in different dimensions as illustrated in Fig. 4:

- Different kinds of image regions: pixels, lines, columns, or rectangles.
- Different region sizes: smaller or larger rectangles and line or column batches.

Depending on the geometry and other properties of the virtual scene, and the different implementation details of the algorithm, the computational load can vary greatly from one region to another [31]. For this reason, we need a safe mechanism to distribute tasks among threads. This assignment can be static (pre-assigned per thread) or dynamic (using a concurrent task queue).

Listing 6 shows the implementation of a static parallelization assignment, distinguishing between the generation of the regions and the rendering process. In this strategy, prefixed indices are computed as a vector of regions and threads access to such indices during the rendering process.

Since it is impossible to estimate the computational load of each task beforehand, a dynamic assignment is likely to be more efficient. Listing 7 shows this implementation, where the vector of regions is replaced by the `ConcurrentBoundedQueue` class from the former lab, including the enqueue and dequeue operations in the generation and rendering processes, respectively. Fig. 5 depicts a diagram of such a concurrent task queue where a main thread generates and enqueues tasks, whereas multiple worker threads dequeue tasks and perform the rendering process in parallel. Fi-

⁴ Temperature can be found in the path `/sys/devices/virtual/thermal/thermal_zone0/temp`.

⁵ <http://www.kevinbeason.com/smallpt/>.

```

struct Region {
    int row0, col0, row1, col1, spp;
};

//Producer
void generate() {
    Image image(width,height);
    std::vector<Region> regions;
    for (region in <regions according to strategy>)
        regions.push_back(region);
    std::vector<std::thread> threads;
    for (int i=0; i<n_th; ++i)
        threads.push_back(std::thread(render, i, n_th, regions,
image));
    for (auto &t : threads)
        t.join();
}

//Consumer
void render(unsigned int thread, unsigned int nthreads,
std::vector<Region>& parts, Image& image) {
    for (int p=thread; p<parts.size(); p+=nthreads)
        for (int row=parts[p].row0; row<parts[p].row1; ++row)
            for (int col=parts[p].col0; col<parts[p].col1; ++col)
                for (int s=0; s<spp; ++s) //rays per pixel
                    image(row, col) += calculatepixel(row, col);
}

```

Listing 6: Static parallelization strategy for computer graphics algorithms.

```

//Producer (main thread)
void generate() {
    Image image(width,height);
    std::vector<std::thread> threads;
    ConcurrentBoundedQueue<Region> regions;
    for (int i=0; i<n_th; ++i)
        threads.push_back(std::thread(render, i, n_th, regions,
image));
    for (region in <regions according to strategy>)
        regions.enqueue(region);
    for (auto &t : threads)
        t.join();
}

//Consumer (worker threads)
void render(unsigned int thread, unsigned int nthreads,
ConcurrentBoundedQueue<Region>& parts, Image& image) {
    while (!parts.done()) {
        Region part = parts.dequeue();
        for (int row=part.row0; row<part.row1; ++row)
            for (int col=part.col0; col<part.col1; ++col)
                for (int s=0; s<spp; ++s) //rays per pixel
                    image(row, col) += calculatepixel(row, col);
    }
}

```

Listing 7: Dynamic parallelization strategy for computer graphics algorithms.




Fig. 5. Concurrent thread-safe task queue to assign tasks to different worker threads.

nally, notice too that both static and dynamic strategies are orthogonal to the region distributions depicted in Fig. 4.

The students should identify the different pros and cons of each of the approaches, analyzing and justifying their impact on performance. For instance, the students should answer questions such as *which is the optimal region size? Which of the mutex approaches of the task queue work best and under which circumstances?*

The final part of the lab notes that the proposed approach relies on *low-level* programming constructs that are helpful to showcase the interactions with the rest of the abstraction levels. However, to boost productivity and get the most of heterogeneous systems, which are standard nowadays, students should be advised to opt for higher level approaches [15,19,34].

Overall, the implementation and parallelization of the path-tracing algorithm together with the performance evaluation of each mutex will help students understand and analyze the effect of low-level mechanisms, decisions, and implementation details with high-level applications and algorithms, which will reinforce the integrated view of a computer system.

4. Experimental environment

To consolidate the overall view of the presented computer system, we propose to use the same single-board computer in all the labs. To this end, we have analyzed a subset of commonly used boards that fulfilled two key restrictions: low-cost (price below 50 \$) and multiprocessing (parallelism experiments cannot be run in single-core boards).

The selected boards are Raspberry Pi 3 Model B [40], ClockworkPi,⁶ Rock64,⁷ AML-S905X-CC (Le Potato),⁸ Orange Pi Zero Plus [29], NanoPi M1 Plus,⁹ and Pine A64-LTS.¹⁰

Table 2 summarizes the most relevant hardware and software requirements for the development of this experience, and which of them are met by the selected boards. Of course, there are other boards offering better performance or more functionality but at a higher cost such as DragonBoard 410C [33], HiKey 960,¹¹ or BeagleBoard X-15 [8]. Such highly-priced boards are not considered in this study.

The list of requirements is mainly focused on the subset of courses taking part in the presented experience. Nevertheless, it is desirable to choose a base board that allows future expansions by adding more courses to the experience. Therefore, we consider a broader range of requirements that would facilitate the use of the selected board for additional courses, such as Computer Architecture and Organization, Systems Administration, Computer Networks, Security, Artificial Intelligence, Machine Learning, Embedded Systems, Robotics, Video-games, or Computer Vision, among others.

Considering the results from our study of boards, requirements, and potential courses that could use them, Raspberry Pi, Orange Pi, and NanoPi turn out good choices to be used in our experience, since they meet all the requirements but the JTAG support. However, we finally chose Raspberry Pi primarily due to its broader usage and large amounts of open source and available materials [41].

5. Technical results

This section presents the main technical results and conclusions that should be obtained by students from the proposed lab assignments. More precisely, the impact on the system performance and chip temperature obtained for every mutex implementation is analyzed under different contention scenarios.

All the experiments are run in a Raspberry Pi 3 Model B, which includes a quad-core processor where each core is single-thread.

⁶ https://wiki.clockworkpi.com/index.php/Main_Page.

⁷ <https://wiki.pine64.org/index.php?title=ROCK64>.

⁸ <https://libre.computer/products/boards/aml-s905x-cc/>.

⁹ http://wiki.friendlyarm.com/wiki/index.php/NanoPi_M1_Plus.

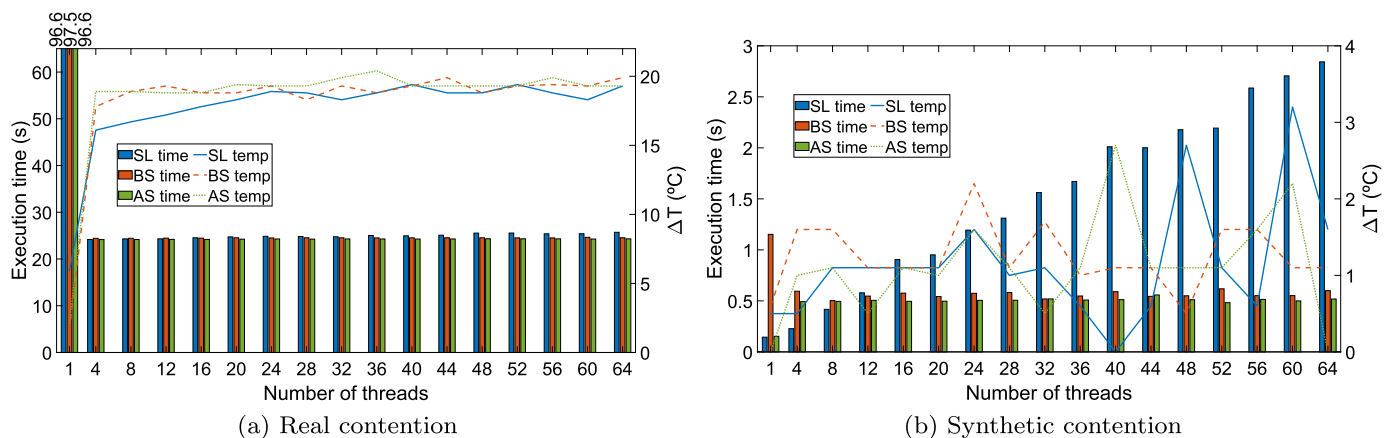
¹⁰ https://wiki.pine64.org/index.php?title=Main_Page.

¹¹ <https://www.96boards.org/documentation/consumer/hikey960/hardware-docs/hardware-user-manual.md.html>.

Table 2

Hardware (H) and Software (S) requirements evaluated for the following boards: Raspberry Pi 3 Model B (RP), ClockworkPi (CP), Rock64 (RC), Le Potato (LP), Orange Pi zero plus (OP), NanoPi M1 Plus (NP), and Pine A64-LTS (PA).

Type	Description	RP	CP	RC	LP	OP	NP	PA
H	JTAG	X	X	X	X	X	X	X
H	Ethernet	✓	✓	✓	✓	✓	✓	✓
H	WiFi	✓	✓	X	X	✓	✓	X
H	Camera	✓	X	X	X	✓	✓	X
H	Virtualization support	✓	✓	✓	✓	✓	✓	✓
H	I/O Extensions (screen, buttons...)	✓	✓	✓	✓	✓	✓	✓
H	GPU	✓	✓	✓	✓	✓	✓	✓
S	Development Framework options	✓	X	✓	X	✓	✓	✓
S	SDK and runtime GPU support	✓	X	✓	X	✓	✓	✓
S	High-level/Standard OS support	✓	X	✓	✓	✓	✓	✓
H&S	Bare metal (no OS) support	✓	✓	✓	✓	✓	✓	✓

**Fig. 6.** Performance and temperature of the OS mutex implementations varying the number of threads.

We assumed a fixed CPU frequency of 1.4 GHz (*performance governor*) for all the experiments to guarantee reproducibility. The OS is an Ubuntu 18.04.3 LTS release with a gcc 7.5.0 compiler. Before running each experiment, we wait until the chip temperature decreases to a defined threshold of 60 °C (the reported experiments were carried out during summer and the board does not include any heat dissipator nor fans). Once an experiment finishes, we calculate the execution time and the chip temperature increase and wait for the CPU to cool down before running the next experiment.

All the presented experiments are run from 1 to 64 threads. Taking into account the quad-core processor and the fact that the studied applications are CPU-bound, the thread oversubscription (i.e., a number of threads higher than the number of physical cores) penalizes performance, which is a key insight for students. In addition, another important key insight from thread oversubscription is that the implementation of each mutex affects performance differently.

5.1. OS level

The first lab assignment deals with the implementation of the concurrent bounded queues and the experiments are limited to ensure the correctness of such queues (Section 3.2.1). In the second lab, the first assignment refers to the implementation of the SL, BS, and AS mutexes using futexes (Section 3.2.2). This section evaluates such implementations.

Bars and lines in Fig. 6 show respectively the execution time in seconds (s) and the chip temperature increase in Celsius degrees (°C). Threads access a shared variable protected with a mutex. In particular, each thread acquiring the lock increments by one the value of this variable and then releases the lock. The execution finishes when the shared variable reaches a value of a million.

Two different contention scenarios are considered, referred to as real and synthetic. In the real scenario, a thread releasing the lock computes some private work consisting of a series of trigonometric functions with the shared variable as input. On the other hand, in the synthetic scenario, after releasing the lock, a thread immediately competes for the lock without performing any private work. The latter scenario covers an extreme case where students observe how a change in the amount of private work leads to unexpected conclusions.

In the real scenario, all the mutexes obtain a very similar performance for a given number of threads greater or equal than 4. For a low number of threads between 1 and 12, the overhead of always invoking a futex system call in the unlock method and the subsequent context switch in BS leads to a slight increase of the execution time compared to SL and AS. However, as the number of threads increases, SL progressively enlarges a bit the execution, since threads spin in a more disputed lock. On the other hand, the sleep mutexes maintain the same execution time as the number of threads increases. With this experiment students realize that, for this kind of application, the thread oversubscription does not improve performance but, on the contrary, depending on the mutex implementation, performance can be hurt.

The increase in the chip temperature is quite steady for a number of threads greater than one, and both the private work and the chip temperature limit (around 80 °C) prevents from obtaining significant CPU temperature differences among the different mutexes. However, SL reaches a lower temperature in most cases, which suggests that invoking futex systems calls has a greater thermal signature.

There is no parallelism to be exploited in the synthetic scenario, meaning that single-threaded executions should exhibit the best

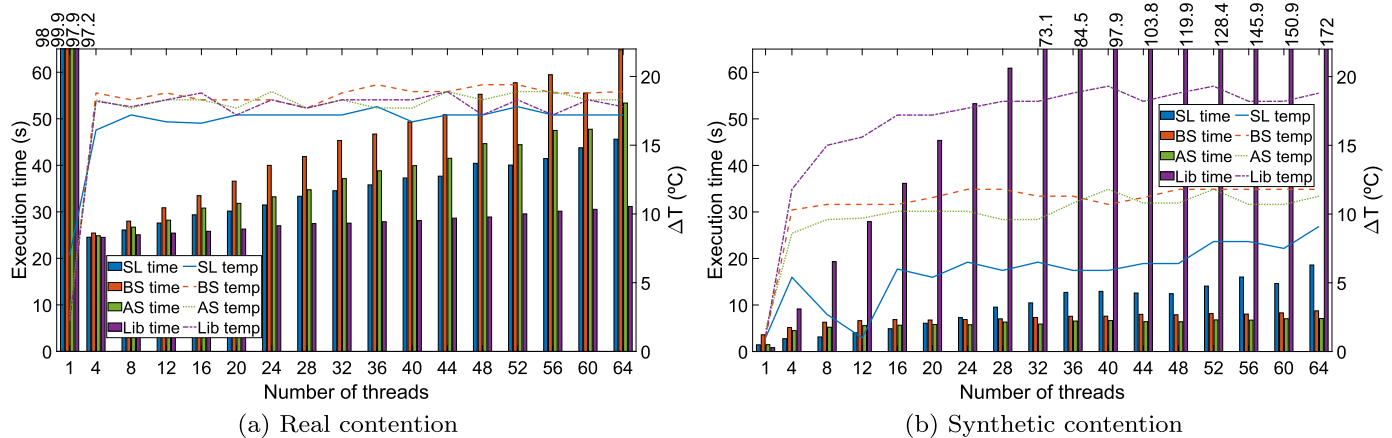


Fig. 7. Performance and temperature of the OS mutex implementations protecting a concurrent bounded queue. A library mutex is included for comparison purposes.

performance for a given mutex implementation. This is the case for both SL and AS mutexes. For the studied mutexes, SL obtains the lowest execution time both in the single-threaded execution and when the number of threads coincides with the number of physical cores. This confirms the overhead of the futex system calls in both BS and AS approaches. On the other hand, similarly to the previous scenario, under the thread oversubscription, spinning largely increases the execution time over the sleep approaches, whereas such mutexes maintain roughly the same performance.

An interesting observation is that, for a single thread, BS has a much larger execution time not only compared to the other mutexes but also compared to itself when multiple threads are considered. This confirms the overhead of always invoking costly `futex_wake` system calls in the unlock function even when there are no waiter threads in the lock. Regarding the chip temperature, the execution time is not sufficiently large to observe significant temperature differences before and after the execution.

5.2. Library level

This section refers to the evaluation of the OS second lab assignment, where the previous mutexes are used to support the concurrent task queue. Fig. 7 shows the results. A library mutex from the DCSP lab is included for comparison purposes. The length of the queue is a million of elements and its initial state is full. Threads within the critical section dequeue values from the queue until it is empty. The private work of each thread is the same as described in the previous section.

The real scenario confirms that the best performing number of threads is 4, where both SL and the library mutex obtain the best performance. This suggests that the library mutex implements a spin-lock mechanism for such a number of threads. The thread oversubscription leads to larger execution times, especially for SL, BS, and AS. The library mutex clearly outperforms all the other mutexes for any number of threads greater than 4, suggesting that, differently from the spin-lock and OS-based solutions, the library implementation possibly throttles unnecessary threads according to the number of processor cores and/or exploits an internal userspace sleep-queue to minimize thread context switches. Contrary to the previous experiments, SL always performs better than AS and BS mutexes. This is mainly due to the additional overhead of the futex system calls from the new suspend and wakeup methods in the Semaphore class. As expected, AS largely reduces the execution time with respect to BS due to the additional cost of always invoking `futex_wake` system calls. Similarly to the previous study, SL shows a slightly lower increase in the CPU temperature compared to the other mutexes, including the library mutex.

Surprisingly, the synthetic scenario greatly affects the library mutex with the number of threads, which points out that such an extreme contention scenario has not been considered in the library mutex development. The higher contention also affects SL, although in a lesser way. In this scenario, with a sufficient number of threads, the overhead of the system calls and subsequent context switches compensates the distribution of CPU time among all the active threads spinning in the lock. Removing the private work allows to see significant temperature differences. The results confirm the lower thermal signature of SL, even in those cases where SL enlarges the execution time with respect to the other approaches. On the other hand, the library mutex shows the highest temperature increase, most likely due to the extended execution times.

Overall, the library mutex is a convenient choice in scenarios where there is a relatively high amount of private work to be done. In this case, the library likely exploits a hybrid management by combining spin-lock, thread throttling, and/or a sleep-queue to provide an adaptive mechanism to the most frequent case. However, in the synthetic scenario, the library heuristics fail to adapt to such an atypical case, making the OS-based mutexes the preferable choice.

5.3. Assembly level

This study evaluates the SL-ASM and LPS-ASM mutexes protecting the concurrent task queue (see Section 3.2.3). The initial state of the queue and the amount of private work is the same as in the previous study. Fig. 8 depicts the results. For illustrative purposes, just the best performing mutexes with 4 threads from the previous study (library and SL for real and synthetic contention, respectively) are shown.

Similarly to the previous studies, the spin-lock solution increases the execution time with the thread oversubscription in the real contention scenario. Compared to SL-ASM, LPS-ASM has a lower execution time since putting cores in a low-power state until a context switch is triggered reduces the lock contention. However, this does not prevent the library mutex to obtain better performance thanks to an enhanced management of the lock. Differently from previous approaches, the LPS-ASM alternative shows a temperature reduction with the thread oversubscription. This is mainly due to, with a higher number of threads, the chance to put cores in a low-power state increases. On the contrary, the steady temperature of the library mutex suggests that such an implementation does not change the state of the cores.

In the synthetic scenario, both spin-lock alternatives progressively increase the execution time with the number of threads.

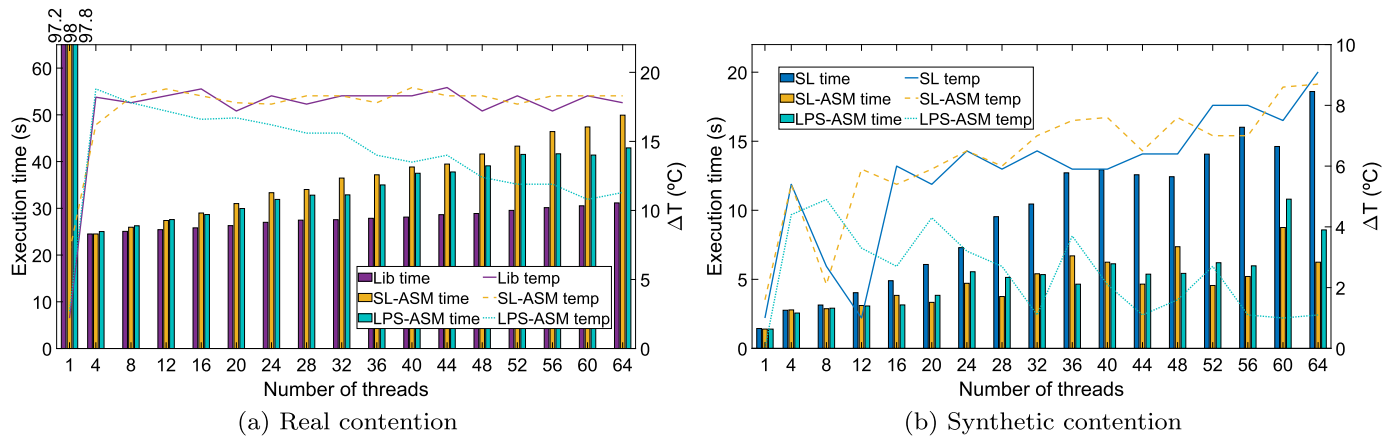


Fig. 8. Performance and temperature of assembly mutexes protecting a concurrent bounded queue. The best performing solutions from the previous study are also shown.

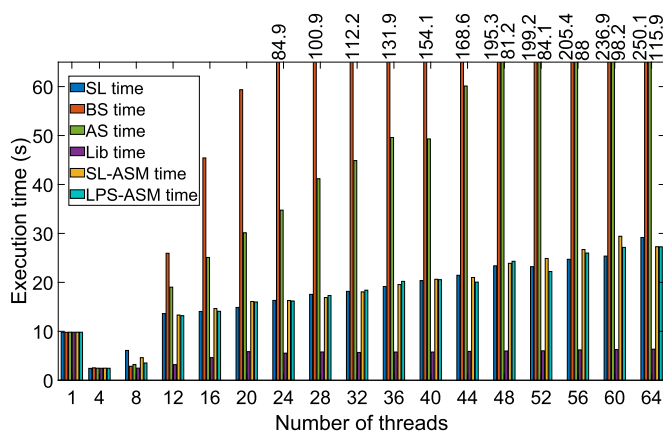


Fig. 9. Performance of all the mutex implementations at the path-tracing application level assuming a dynamic parallelization strategy.

However, implementing the spin-lock in a higher abstraction level with respect to the ISA level introduces a performance overhead according to the timing differences between SL and SL-ASM. Finally, the LPS-ASM solution clearly reduces the temperature with respect to both spin-lock approaches, which show very similar temperature numbers.

5.4. Application level

This section evaluates all the mutexes in the ray tracing application using the dynamic parallelization strategy with the concurrent task queue (see Section 3.2.4). Fig. 9 plots the results, which are restricted to performance for illustrative purposes. The rendered scene in all the experiments is *forest*.¹² For simplification purposes, the image is partitioned in fixed-size rectangular regions.

As expected from the real scenarios of the previous studies, the best performing number of threads is 4 according to the number of physical cores, where all the analyzed mutexes exhibit a very similar performance. In addition, the performance of the ray tracer scales with the number of threads, since the execution time with 4 threads reduces by 4× with respect to the single-threaded performance. Results under the thread oversubscription corroborate the previous findings from lower abstraction levels: the OS overhead largely penalizes the AS mutex and especially the BS mutex, whereas the OS-free solutions like SL and the assembly mutexes

show similar results. Finally, the library mutex is scarcely affected by the thread oversubscription.

6. Experience assessment

This section provides a qualitative assessment of all the proposed labs. The attendance to such labs was voluntary (37 students participated in each lab of the DCSP and OS courses, and 12 students participated in each lab of the MP and CG courses). All labs were scheduled after the completion of the courses, giving the students the opportunity to compare between the current lab assignments (i.e., no direct interactions with any other lab) and the proposed lab assignments. Notice too that this experience has been carried out for three consecutive academic years, meaning that those students that have participated in the latest CG lab have also participated in the remaining ones. Assessment results collected for the DCSP and OS labs correspond to two years: current and previous, whereas results for the MP and CG labs correspond to the previous and current year, respectively. In the first year, the first two labs were run without collecting assessment results.

Two different surveys were designed for each lab, referred to as pre-survey and post-survey. Students filled out the surveys before and after completing the proposed lab sessions. Every pre-survey consists of two types of questions: a series of questions for measuring the perception of students about the interactions among the courses of the CE program in general and particularly among the involved courses in the experience, and 22 theoretical and practical questions for assessing the knowledge of students about the covered PDC topics.

Every post-survey is structured in four parts: a set of questions for stressing in the interactions among the courses, a series of questions for valuing the proposed lab experiences, other questions referring to the obtained technical results, and the same 22 theoretical and practical questions from the pre-surveys. The main conclusions extracted from these surveys are summarized next.

Before the lab sessions, the students considered that all the courses of the program are somehow related (8.0 points out of 10, where the edge scores 0 and 10 indicate no relation at all and totally related, respectively), but that the faculty should make an effort to make this relation more explicit (the effort was rated 7.9 out of 10). Regarding the courses involved in this learning experience, the students were very conclusive. All of them considered that the four courses are strongly related to other courses of the program, including between themselves. The responses of the post-surveys strengthened the previous results, since the students conveyed that the courses are more related to each other (8.7 out of 10).

¹² <http://www.kevinbeason.com/smallpt/extraScenes.txt>.

Table 3

Correct answer rates of the common questions for pre and post surveys.

DCSP						
	Q1	Q2	Q3	Q4		
Pre	100	72.9	27.1	56.7		
Post	100	97.3	43.2	59.5		
OS						
Q5	Q6	Q7	Q8	Q9	Q10	
97.3	81.1	48.7	67.6	81.1	10.8	
97.3	97.3	91.9	100	100	100	
MP						
Q11	Q12	Q13	Q14	Q15		
75	41.7	41.7	41.7	33.3		
100	91.7	83.3	100	75		
CG						
Q16	Q17	Q18	Q19	Q20	Q21	Q22
75.0	83.3	33.3	41.7	33.3	50.0	58.8
91.7	100	66.7	83.3	100	83.3	83.3

Table 3 compares the correct answer rates of the 22 theoretical and practical questions that were included in both kind of surveys. Questions Q1-Q4 were proposed by the faculty of the DCSP course and mainly focus on the concept of mutual exclusion, the semantics of semaphores, and the use of semaphores to solve certain critical section problems. Questions Q5-Q10 refer to the OS session and are about the low-level requirements to implement different types of semaphores. Questions Q11-Q15 refer to the MP course and are about the role of atomicity and consistency in thread synchronization. Questions Q16-Q22 correspond to the CG course and deal with the impact of the different abstraction levels in the parallel ray tracing application. The reader is referred to Appendix A for further details about these questions, including the possible and the correct/desired answers, highlighted in boldface.

In general, the post-survey rates improve the correct answer rates obtained in the pre-surveys. Overall, the results support the interest of the experience and the usefulness from the students point of view. On average, the correct answer rate of the post-surveys is 88.0% versus 57.3% obtained before the beginning of the sessions. One remarkable aspect to be noticed is the rather low rate of correct answers to questions Q3 and Q4. In our opinion, the main reason is that the conceptual study of the different semantics of the semaphore abstract data type has been presented and discussed two months before the experience has been carried out. For a set of reasons, we are currently thinking about the possibility of changing the order in which concepts are introduced at the OS course, so as to be able to study the semaphore abstraction and implementation closer in time. Maybe this fact will help us to get more information to understand the reason for such a low rate.

When comparing year-on-year results of the surveys, we have shown some striking facts. This is the case of questions Q16 and Q17 of the CG lab, which are the same as questions Q8 and Q9 of the OS lab. In the survey after the OS lab, 100% of the students answered them correctly. However, two years later in the previous survey of the CG lab, these same questions were answered correctly by only 75% and 83.3% of the students, respectively. These results demonstrate the effect of time on students' memory. Fortunately, the CG lab session helps to refresh students' memory by reaching 91.7% and 100% on the post-survey.

The post-surveys revealed that all the lab sessions were well received. All the students completed the lab assignments, and they gave an overall score of 8.8 out of 10 to the quality of the lab designs, the materials and resources, and the faculty assistance. When asked about their opinion in the CG lab, one student men-

tioned that he/she "liked a lot the idea of unifying several courses in a single learning experience, and this should be done more often to consolidate the learning". Another student pointed out that "this is a very positive initiative to put together everything we have seen in multiple courses".

In general, after completing the labs, the students have reached a broader view of the interactions among operating systems, computer architecture, and parallel and distributed computing. As learning outcomes, students discerned among the different mutex implementation alternatives and clearly identified the programmability, execution time, and efficiency trade-offs at each abstraction level.

7. Conclusions

The current structure of the Computer Engineering (CE) program, arranged in isolated courses, causes students to lose sight of the overall view of a typical computer system organized in abstraction levels. This paper has presented a learning experience that aims to reinforce this vision as a whole.

The presented experience covers the abstraction levels of Application, Library, Operating System, and Instruction Set Architecture, and consists in the implementation of a parallel ray tracing algorithm that uses a concurrent queue to assign tasks to different execution threads. The accesses in mutual exclusion to this queue are managed by mutexes implemented with either library functions, system calls, or assembly instructions.

The aforementioned abstraction levels have been introduced and related to each other in a subset of laboratories from different courses of a CE program, allowing students to consolidate the concepts of parallelism, concurrency, atomicity, and consistency. This paper has presented the structure and contents of each proposed lab, as well as the interactions with the remaining labs. In addition, a detailed study of the hardware and software requirements and the consequent choice of Raspberry Pi as the common hardware development platform is also discussed.

Experimental results consisted of a technical evaluation and an assessment study of the proposed learning experience. The technical results referred to an evaluation and discussion of the performance and temperature differences of the implemented mutexes in each lab. The experience assessment consisted of a series of pre/post surveys. Most students pointed out an enhancement in the design of the labs and a greater exposure to the relations between courses. In addition, the students showed an enhancement in the integrated perception of the addressed concepts and the acquisition of the knowledge, since the correct responses to the technical questions from the surveys improved by 30.7% after the experience.

Declaration of competing interest

There are no conflicts of interest to report for this submission.

Acknowledgments

All authors acknowledge support from grants (1) PIIDUZ_18_246 from *Universidad de Zaragoza*, (2) PID2019-105660RB-C21 / AEI / 10.13039/501100011033 from *Agencia Estatal de Investigación* (AEI) and European Regional Development Fund (ERDF), (3) gaZ: T58_20R research group from Aragon Government and European Social Fund (ESF), (4) 2014-2020 "Construyendo Europa desde Aragón" from ERDF, (5) TIN2017-84796-C2-2-R from Spanish *Ministerio de Economía y Competitividad*, (6) DisCo-T21-20R from Aragon Government, (7) PID2019-105004GB-I00 from Spanish Ministry of Science and Innovation, and (8) CHAMELEON, grant No. 682080 from European Research Council (ERC) under the EU's Horizon 2020 research and innovation programme.

Appendix A. Surveys

A.1. DCSP survey

Q1 Do you understand the importance of mutual exclusion?

- **Yes**
- No
- I have a rough idea

Q2 When a P process is said to use an S semaphore with busy-wait semantics, which of the following operations refers to wait (S) ?:

- Option 1:

```
< if S.V>0
  S.V := S.V-1
else
  S.L := S.L ∪ {P}
  P.state := blocked
end
>
```

- **Option 2:**

```
< await S.V>0
  S.V := S.V-1
>
```

- To anyone, both are equivalent

Q3 If the semantics for semaphores in the first two answers from the previous question are taken into account, which of the following statements can be considered as true?

- Since they are equivalent, there are no differences in terms of liveness and safety properties
- **There are no differences in terms of safety properties, but there are differences in terms of liveness**
- There are no differences in terms of liveness properties, but there are differences in terms of safety
- Although they are not equivalent, there are no differences in terms of liveness and safety properties

Q4 Consider the following scheme to solve the problem of the critical section, S being a semaphore with busy-wait semantics. Which of the following statements is correct?

Semaphore S:=1	Process P	Process Q
	loop forever	loop forever
	SNC	SNC
	wait (S)	wait (S)
	SC	SC
	signal (S)	signal (S)
	end	end
	end	end

- **The solution can generate fairness problems**
- There are no fairness issues
- There will be fairness issues depending on the first process that gets access to the semaphore

A.2. OS survey

Q5 Do you know the advantages and shortcomings of a mutex with/without active waiting?

- **Yes**
- No
- There are no differences

Q6 Would you be able to implement a mutex with busy waiting?

- **Yes**
- No
- I am not sure

Q7 Would you be able to implement a mutex without busy waiting?

- **Yes**

- No
- I am not sure

Q8 What do you consider essential to implement a mutex with busy waiting (multiple choices can be selected)?

- **Atomic memory reading and writing instructions**
- OS support
- **A shared memory space**
- Nothing, any system supports it by default

Q9 What do you consider essential to implement a mutex without busy waiting (multiple choices can be selected)?

- **Atomic memory reading and writing instructions**
- **OS support**
- **A shared memory space**
- Nothing, any system supports it by default

Q10 Do you know what a futex is?

- **Yes**
- No
- I have a rough idea

A.3. MP survey

Q11 The consistency model defined by ARMv8 is...

- Sequential
- **Relaxed**
- None of them

Q12 Which of the following instructions puts the processor on a low power state?

- sevl
- dbm
- **wfe**
- All of the above
- None of the above

Q13 For a *fetch and add*, would you use any of the following ARMv8 instructions?

- wfe, sevl
- DMB, DSB, ISB
- **ldaxr, stlr**
- None of the above

Q14 What of the following synchronization alternatives would you use in a low-contention scenario for the next program?

```
// run concurrently by several threads
long long my_add = 0;
while (work_index < n_elements) {

    // how do you protect the next code lines?
    int my_work_index_ini = work_index;
    work_index = work_index + chunk;
    int my_work_index_end = work_index;

    for (int i=my_work_index_ini; i<my_work_index_end; i++)
        my_add = my_add + v_elems[i];
}
```

- **Fetch and add**
- Mutex
- Energy-efficient mutex

Q15 What of the following synchronization alternatives would you use in a high-contention scenario for the same program as in the previous question?

- *Fetch and add*
- Mutex
- **Energy-efficient mutex**

A.4. CG survey

Q16 What are the minimal requirements to implement a non busy-wait mutex (several choices can be valid)?

- **Atomic instructions**

- **OS support**
 - **A shared memory space**
 - Nothing, all systems support non busy-wait mutexes
- Q17 What are the minimal requirements to implement a busy-wait mutex (several choices can be valid)?
- **Atomic instructions**
 - OS support
 - **A shared memory space**
 - Nothing, all systems support non busy-wait mutexes
- Q18 For a parallel application such as a ray tracer, which implementation would be faster in a scenario where the number of logical threads is *lower* than the number of physical execution contexts of the processor?
- OS mutex
 - Library mutex
 - Assembly mutex
 - **All three should have a similar behavior**
- Q19 For a parallel application such as a ray tracer, which implementation would be faster in a scenario where the number of logical threads is *equal* to the number of physical execution contexts of the processor?
- OS mutex
 - Library mutex
 - Assembly mutex
 - **All three should have a similar behavior**
- Q20 For a parallel application such as a ray tracer, which implementation would be faster in a scenario where the number of logical threads is *higher* than the number of physical execution contexts of the processor?
- OS mutex
 - **Library mutex**
 - Assembly mutex
 - All three should have a similar behavior
- Q21 Many algorithms, such as ray tracing, whose output is a matrix of pixels (or an image), are easily parallelizable. Which is the best parallelization strategy for such algorithms?
- Static: pixels are pre-assigned to a worker thread by rows, columns, or regions. In this approach, workers perform their tasks without synchronizing with other workers
 - Dynamic-per-pixel: each worker continuously fetches single-pixel tasks until completion. It requires synchronization among workers
 - **Dynamic-per-region: each worker continuously fetches region tasks until completion. It requires synchronization among workers**
 - Hard to guess: the best strategy depends on the variability of the execution time of each task
- Q22 In a parallel ray tracing algorithm, where each pixel can be independently generated, assuming a dynamic-per-region strategy, under which circumstance reducing the region size can improve performance?
- When there is variability in the pixel-generation time
 - When there are image areas where is known beforehand that the pixel-generation time is smaller than that of other areas
 - **When the region task delay is significantly larger than the synchronization delay produced by the region**
 - When the resulting image has many color disturbances

References

- [1] ACM/IEEE, Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science, 2013.
- [2] ACM/IEEE, Computer Engineering Curricula 2016: Curriculum Guidelines for Undergraduate Degree Programs in Computer Engineering, 2016.
- [3] G.R. Andrews, Concurrent Programming. Principles and Practice, 1st edition, The Benjamin/Cummings Publishing Company, Inc., 1991.
- [4] ARM, ARM DS-5 Development Studio Examples, 2018.
- [5] R.H. Arpaci-Dusseau, A.C. Arpaci-Dusseau, Operating Systems: Three Easy Pieces, 1st edition, Arpaci-Dusseau Books, LLC, 2018.
- [6] S.A. Bogaerts, One step at a time: parallelism in an introductory programming course, J. Parallel Distrib. Comput. 105 (2017) 4–17.
- [7] M. Burtscher, W. Peng, A. Qasem, H. Shi, D. Tamir, H. Thiry, A module-based approach to adopting the 2013 ACM curricular recommendations on parallel computing, in: Proceedings of the 46th ACM Technical Symposium on Computer Science Education, 2015, pp. 36–41.
- [8] G. Coley, BeagleBoard X15 system reference manual, BeagleBoard.org, 2016.
- [9] J. Cuenca, D. Giménez, A parallel programming course based on an execution time-energy consumption optimization problem, in: Proceedings of the IEEE International Parallel and Distributed Processing Symposium Workshops, 2016, pp. 996–1003.
- [10] U. Drepper, Futexes are tricky, <http://people.redhat.com/drepper/futex.pdf>, 2011.
- [11] U. Drepper, I. Molnar, The native POSIX thread library for Linux, <https://akkadia.org/drepper/nptl-design.pdf>, 2005.
- [12] J. Eckroth, A course on big data analytics, J. Parallel Distrib. Comput. 118 (P1) (2018) 166–176.
- [13] C. Ferner, B. Wilkinson, B. Heath, Toward using higher-level abstractions to teach parallel computing, in: Proceedings of the IEEE International Parallel and Distributed Processing Symposium Workshops, 2013, pp. 1291–1296.
- [14] H. Franke, R. Russell, M. Kirkwood, Fuss, futexes and furwocks: fast user-level locking in Linux, in: Proceedings of the Ottawa Linux Symposium, 2002, pp. 479–495.
- [15] B. Gaster, L. Howes, D.R. Kaeli, P. Mistry, D. Schaa, Heterogeneous Computing with OpenCL, 1st edition, Morgan Kaufmann Publishers Inc., 2011.
- [16] N. Giacaman, O. Sinnen EA, Research-infused teaching of parallel programming concepts for undergraduate software engineering students, in: IEEE International Parallel Distributed Processing Symposium Workshops, 2014, pp. 1099–1105.
- [17] D. Ginat, Y. Blau, Multiple levels of abstraction in algorithmic problem solving, in: Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education, 2017, pp. 237–242.
- [18] M. Grossman, M. Aziz, H. Chi, A. Tibrewal, S. Imam, V. Sarkar, Pedagogy and tools for teaching parallel computing at the sophomore undergraduate level, J. Parallel Distrib. Comput. 105 (2017) 18–30.
- [19] M. Haidl, S. Gorlatch, PACXX: towards a unified programming model for programming accelerators using C++14, in: Proceedings of the LLVM Compiler Infrastructure in HPC, 2014, pp. 1–11.
- [20] D.J. John, S.J. Thomas, Parallel and distributed computing across the computer science curriculum, in: Proceedings of the IEEE International Parallel and Distributed Processing Symposium Workshops, 2014, pp. 1085–1090.
- [21] J. Kramer, Is abstraction the key to computing?, Commun. ACM 50 (4) (2007) 36–42.
- [22] S. Kumar, Research-oriented teaching of PDC topics in integration with other undergraduate courses at multiple levels: a multi-year report, J. Parallel Distrib. Comput. 105 (2017) 92–104.
- [23] V. Kumar, Introduction to Parallel Computing, 2nd edition, Addison-Wesley Longman Publishing Co., Inc., 2002.
- [24] B. Levandowski, D. Perouli, D. Brylow, Using embedded Xinu and the raspberry Pi 3 to teach parallel computing in assembly programming, in: Proceedings of the IEEE International Parallel and Distributed Processing Symposium Workshops, 2019, pp. 334–341.
- [25] C. Lupo, Z.J. Wood, C. Victorino, Cross teaching parallelism and ray tracing: a project-based approach to teaching applied parallel computing, in: Proceedings of the 43rd ACM Technical Symposium on Computer Science Education, 2012, pp. 523–528.
- [26] S.J. Matthews, Teaching with parallella: a first look in an undergraduate parallel computing course, J. Comput. Sci. Coll. 31 (3) (2016) 18–27.
- [27] S.J. Matthews, J.C. Adams, R.A. Brown, E. Shoop, Portable parallel computing with the raspberry Pi, in: Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education, 2018, pp. 92–97.
- [28] T. Newhall, A. Danner, K.C. Webb, Pervasive parallel and distributed computing in a liberal arts college curriculum, J. Parallel Distrib. Comput. 105 (2017) 53–62.
- [29] Orange Pi Zero H2 User Manual, version 0.9.1, Shenzhen Xunlong Software Co., Ltd.,
- [30] S. Petit, J. Sahuquillo, M.E. Gómez, V. Selfa, A research-oriented course on advanced multicore architecture: contents and active learning methodologies, J. Parallel Distrib. Comput. 105 (2017) 63–72.
- [31] M. Pharr, W. Jakob, G. Humphreys, Physically Based Rendering: From Theory to Implementation, 3rd edition, Morgan Kaufmann Publishers Inc., 2017.
- [32] S.K. Prasad, A. Chtchelkanova, F. Dehne, M. Gouda, A. Gupta, J. Jaja, K. Kant, A. La Salle, R. LeBlanc, A. Lumsdaine, D. Padua, M. Parashar, V. Prasanna, Y. Robert, A. Rosenberg, S. Sahni, B. Shirazi, A. Sussman, C. Weems, J. Wu, NSF/IEEE-TCPP curriculum initiative on parallel and distributed computing - core topics for undergraduates, version I, <http://tcpp.cs.gsu.edu/curriculum/>, 2012.

- [33] Qualcomm, DragonBoard™410C based on Qualcomm®Snapdragon™410E processor. Peripherals Programming Guide Linux Android, Qualcomm Technologies, Inc., 2016.
- [34] J. Reinders, B. Ashbaugh, J. Brodman, M. Kinsner, J. Pennycook, X. Tian, Data Parallel C++: Mastering DPC++ for Programming of Heterogeneous Systems Using C++ and SYCL, Springer Nature, 2021.
- [35] E. Saule, Experiences on teaching parallel and distributed computing for undergraduates, in: IEEE International Parallel and Distributed Processing Symposium Workshops, 2018, pp. 361–368.
- [36] A. Silberschatz, P.B. Galvin, G. Gagne, Operating System Concepts, 9th edition, Wiley Publishing, 2012.
- [37] D.J. Sorin, M.D. Hill, D.A. Wood, A Primer on Memory Consistency and Cache Coherence, 1st edition, Morgan & Claypool Publishers, 2011.
- [38] S. Srivastava, M. Smith, A. Ghimire, S. Gao, Assessing the integration of parallel and distributed computing in early undergraduate computer science curriculum using unplugged activities, in: IEEE/ACM Workshop on Education for High-Performance Computing, 2019, pp. 17–24.
- [39] W. Stallings, Operating Systems: Internals and Design Principles, 6th edition, Prentice Hall Press, 2008.
- [40] E. Upton, G. Halfacree, Raspberry Pi User Guide, John Wiley & Sons Ltd., 2014.
- [41] E. Upton, J. Duntemann, R. Roberts, T. Mamtara, B. Everard, Learning Computer Architecture with Raspberry Pi, 1st edition, Wiley Publishing, 2016.
- [42] L.B.A. Vasconcelos, F.A.L. Soares, P.H.M.M. Penna, M.V. Machado, L.F.W. Góes, C.A.P.S. Martins, H.C. Freitas, Teaching parallel programming to freshmen in an undergraduate computer science program, in: IEEE Frontiers in Education Conference, 2019, pp. 1–8.
- [43] E. Veach, Robust Monte Carlo Methods for Light Transport Simulation, Ph.D. thesis, Stanford University, 1998.
- [44] A. Williams, C++ Concurrency in Action, Manning Publications, 2012.
- [45] A.A. Younis, R. Sunderraman, M. Metzler, A.G. Bourgeois, Case study: using project based learning to develop parallel programming and soft skills, in: IEEE International Parallel and Distributed Processing Symposium Workshops, 2019, pp. 304–311.



Alejandro Valero received the Ph.D. degree in Computer Engineering from the Universitat Politècnica de València, Spain, in 2013. From 2013 to 2015, he was a visiting researcher with Northeastern University, Boston, MA, USA, and the University of Cambridge, UK. Since 2016, he has been a professor with the Department of Computer Science and Systems Engineering, Universidad de Zaragoza, Spain, where he teaches several courses on computer organization and

has been involved in multiple teaching innovation projects. His research interests include GPU architectures, memory hierarchy design, energy efficiency, fault tolerance, and computer architecture education. Prof. Valero is a member of the Aragon Institute of Engineering Research (I3A) and the HiPEAC European NoE.



Rubén Gran-Tejero has been teaching for the last fourteen years at the Department of Computer Science and Systems Engineering of the University of Zaragoza. He has conducted lectures on operating systems, systems architecture, computer architecture, parallel multiprocessors, and high performance computing. He has been involved in several teaching innovation projects as collaborator and principal investigator. His research interests include hard real-time

systems, hardware for reducing worst-case execution time and energy consumption, efficient processor microarchitecture, and effective programming for parallel and heterogeneous systems.



Darío Suárez-Gracia received the Ph.D. degree in Computer Engineering from the Universidad de Zaragoza, Spain, in 2011. From 2012 to 2015, he was with Qualcomm Research Silicon Valley. He is currently an Associate Professor with the Universidad de Zaragoza. His research interests include parallel programming, heterogeneous computing, fault-tolerance and computer architecture education. He is a recipient of the best paper award at the Workshop on

Computer Architecture Education. He is a member of the Aragon Institute of Engineering Research (I3A), the IEEE Computer Society, ACM, and the HiPEAC European NoE.



Emanuel A. Georgescu is an Undergraduate Student of the Telecommunications Engineering Program at the Universidad de Zaragoza. His undergraduate thesis explores how different mutex implementations impact on the performance of a ray tracer application.



Joaquín Ezpeleta received the M.S. degree in Mathematics and the Ph.D. degree in Computer Science from the University of Zaragoza, Spain. He is a Full Professor of the Dept. of Computer Science and Systems Engineering of the University of Zaragoza, where he conducts lectures on formal methods for sequential and concurrent programming as well as service-oriented architectures. His research focuses on the problem of modeling, analysis, and control syn-

thesis for concurrent systems and the application of formal techniques to help in the development of correct distributed systems based on Internet and cloud technologies, as well as on the parallel processing of data and intensive computing problems.



Pedro Álvarez received the Ph.D. degree in Computer Science Engineering from the University of Zaragoza, Zaragoza, Spain, in 2004. He works as Senior Lecturer Professor at this University since 2000. His current research interests focus on the development and application of novel techniques for the analysis of large-scale data repositories. The proposed solutions have been applied in different fields, such as cybersecurity, e-commerce and e-business, or e-health and

sports. During the last years these research results have been transferred to the academic context in order to understand students' behavior and to improve the learning-teaching methods.



Adolfo Muñoz is an Associate Professor at the Department of Computer Science and Systems Engineering of the University of Zaragoza. For the last fifteen years, he has been teaching in diverse courses related to programming, software engineering and computer graphics in varied undergrad and master level engineering and computer science studies. He has participated in several teaching innovation projects, and has applied innovative methodologies throughout the different courses he has participated in. He is recently involved in a new

master in Robotics, Graphics, and Computer Vision.



Luis M. Ramos is an Associate Professor at the Department of Computer Science and Systems Engineering of the University of Zaragoza. He received the Ph.D. degree in Computer Engineering from the University of Zaragoza in 2009. For the last 25 years he has been teaching in diverse courses related to logic design, computer architecture, high performance computing, and operating systems. He has participated in several teaching innovation projects. His research interests include hardware data prefetching and computer architecture education. He received the Best Paper Award at the First JILP Data Prefetching Championship.

He received the Best Paper Award at the First JILP Data Prefetching Championship.



Pablo Ibáñez received the M.S. degree in Computer Science from the Universitat Politècnica de Catalunya, Spain, in 1989, and the Ph.D. degree in Computer Science from the Universidad de Zaragoza, Spain, in 1998. He is an Associate Professor with the Computer Science and Systems Engineering Department, University of Zaragoza. His research interests include processor microarchitecture, memory hierarchy, parallel computer architecture, and high performance computing applications.

performance computing applications.