

Document downloaded from:

<http://hdl.handle.net/10251/168795>

This paper must be cited as:

Castelló, A.; Mayo Gual, R.; Seo, S.; Balaji, P.; Quintana Ortí, ES.; Peña, AJ. (2020). Analysis of threading libraries for high performance computing. IEEE Transactions on Computers. 69(9):1279-1292. <https://doi.org/10.1109/TC.2020.2970706>



The final publication is available at

<https://doi.org/10.1109/TC.2020.2970706>

Copyright Institute of Electrical and Electronics Engineers

Additional Information

© 2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

# Analysis of Threading Libraries for High-Performance Computing

Adrián Castelló

Universitat Jaume I

Castellón de la Plana, Spain

Email: [adcastel@uji.es](mailto:adcastel@uji.es)

Rafael Mayo Gual

Universitat Jaume I

Castellón de la Plana, Spain

Email: [mayo@uji.es](mailto:mayo@uji.es)

Sangmin Seo

Ground X

Seoul, Korea

Email: [seo.sangmin@gmail.com](mailto:seo.sangmin@gmail.com)

Pavan Balaji

Argonne National Laboratory

Lemmont, USA

Email: [balaji@anl.gov](mailto:balaji@anl.gov)

Enrique S. Quintana-Ortí

Universitat Politècnica de València

València, Spain

Email: [quintana@disca.upv.es](mailto:quintana@disca.upv.es)

Antonio J. Peña

Barcelona Supercomputing Center (BSC)

Barcelona, Spain

Email: [antonio.pena@bsc.es](mailto:antonio.pena@bsc.es)

**Abstract**—With the appearance of multi-/many core machines, applications and runtime systems evolved in order to exploit the new on-node concurrency that brought new software paradigms. POSIX threads (Pthreads) was widely-adopted for that purpose and it remains as the most used threading solution in current hardware. Lightweight thread (LWT) libraries emerged offering lighter mechanisms to tackle the massive concurrency that current hardware is offering. In this paper, we analyze in detail the most representative threading libraries including Pthread- and LWT-based solutions. In addition, to examine the suitability of LWTs for different use cases, we develop a set of microbenchmarks consisting of commonly found OpenMP patterns in current parallel codes, and we compare the results using threading libraries and OpenMP implementations. Moreover, we study the semantics offered by threading libraries in order to expose the similarities among different LWT application programming interfaces and their advantages over Pthreads. This study reveals that LWT libraries outperform solutions based on operating system threads in cases where tasks and nested parallelism are required.

**Index Terms**—Lightweight Threads, OpenMP, GLT, POSIX Threads, Programming Models

## 1. Introduction

Operating system (OS) threads, such as POSIX threads (Pthreads) [1], are employed to leverage the computational power of multi-/many-core architectures. These threads are commonly used in two ways: via the Pthreads implementation or by means of threading-based programming models (PMs) that rely on Pthreads, such as OpenMP [2].

In the past few years, the number of cores per processor has increased steadily, reaching impressive counts such as the 260 cores per socket in the Sunway TaihuLight supercomputer [3], which was ranked #1 for the first time in the June 2016 TOP500 list [4]. This trend indicates that exascale systems may well feature a large number of cores. Therefore, future applications will have to accommodate this massive concurrency by deploying a large number of threads and/or tasks in order to extract a significant fraction of the computational power of such hardware.

Current solutions for extracting on-node parallelism are based on OS threads in both low- or high-level libraries. However, performing thread management in the OS in-

creases the cost of these operations (e.g. creation, context-switch, or synchronization). As a consequence, leveraging OS threads for a massive degree of hardware parallelism may be difficult. In response to this problem, dynamic scheduling and lightweight threads (LWTs) (also known as user-level threads, or ULTs) models were first proposed in [5] in order to deal with the required levels of parallelism, offering more efficient management, context switching and synchronization operations. These thread solutions rely on the usage of threads that are managed in the user-space so that the OS is not involved and, hence, the overhead is lower.

To illustrate this, Figure 1 highlights the time spent when creating OS and user-level threads (labeled as OS and ULT, respectively). In this example, one thread is created for each core in a machine with two Intel Xeon E5-2695v4 (2.10 GHz) CPUs and 128 GB of memory. For the OS thread, we have employed the GNU C 6.1 library [6], and an Argobots (07-2018) thread for the ULT case [7]. The time difference is caused by the implication of the OS and by the features of each type of thread.




Figure 1: Cost of creating OS threads and ULTs.

A number of LWT libraries have been implemented for specific OSs, such as Windows Fibers [8] and Solaris Threads [9], for specific hardware such as TiNy-threads [10] for the Cyclops64 cel-

lular architecture, or for network services such as Capriccio [11]. Other solutions emerged to support specific higher-level PMs. This is the case of Converse Threads [12] [13] for Charm++ [14] and Nanos++ LWTs [15] for task parallelism in OmpSs [16]. Moreover, general-purpose solutions have emerged such as GNU Portable Threads [17], StackThreads/PM [18], ProtoThreads [19], MPC [20], MassiveThreads [21], Qthreads [22] and Argobots [7]. Other solutions that abstract LWT facilities include Cilk [23], Intel TBB [24] and Go [25]. In addition, solutions like Stackless Python [26] and Protothreads [19] are more focused on stackless threads.

In spite of their potential performance benefits, none of these LWT software solutions has been significantly adopted to date. The easier code development via directive-based PMs, in combination with the lack of a standard/specification, hinder portability and require a considerable effort to translate code from one PM to another. In order to tackle this situation, a common application programming interface (API), called Generic Lightweight Threads (GLT), was presented in [27]. This API unifies LWT solutions under a unique set of semantics, becoming the first step toward a standard/specification. GLT is currently implemented on top of Qthreads, MassiveThreads, and Argobots. One further step is presented in [28] and [29], where the OpenMP and OmpSs PMs were implemented on top of the GLT API.

In this paper we demonstrate the usability and performance gain of LWT solutions. We analyze several threading solutions from a semantic point of view, identifying their strong and weak points. Moreover, we offer a detailed performance study using the OpenMP PM because of its position as the *de facto* standard parallel programming model for multi/many-core architectures. Our results reveal that the performance of most of the LWT solutions is similar and that these are as efficient as OS threads in some simple scenarios, while outperforming them in many, more complex cases.

In our previous work [30], we compared several LWT solutions and used the OpenMP PM as the baseline; in this paper we expand that work adding Pthreads library to our semantic and functional analysis of threading libraries and in order to highlight the overhead (if any) introduced by the OpenMP implementations. The purpose of this paper is to present the first comparison of threading libraries from a semantic point of view, along with a complete performance evaluation that aims to demonstrate that LWTs are a promising replacement for Pthreads used both as low-level libraries and as the base of high-level PMs.

The contributions of this paper are: (1) an extensive description of the current and most-used threading solutions; (2) an analysis of their APIs; and (3) a performance analysis designed to illustrate the benefits of leveraging LWTs instead of OS threads.

The rest of the paper is organized as follows. Section 2 briefly reviews related work. Section 3 reviews in detail the threading solutions. Section 4 presents an analysis of the different LWT approaches. Section 5 introduces the different parallel patterns that are analyzed. Section 6 provides implementation details on the microbenchmarks we developed

for this paper. Section 7 analyzes the performance of LWT libraries. Section 8 closes the paper with conclusions.

## 2. Related Work

The use of ULTs to increase concurrency while maintaining performance is not a new topic. The concept of LWT was introduced in [5], focusing on fundamentals such as scheduling, synchronization, and local storage. Converse Threads was later presented in [12] as a low-level LWT library. It supports not only ULTs but also stackless threads called Messages. Qthreads was presented in [22] and compared with the Pthreads library by means of a set of microbenchmarks and applications. This solution increases the number of hierarchical levels to three with an intermediate element known as Worker. MassiveThreads was presented in [21]. This work provides a performance comparison among MassiveThreads, Qthreads, Nanos++, Cilk, and Intel TBB on several benchmarks. Argobots was presented in [7] with microbenchmark and application evaluations against Qthreads and MassiveThreads. This library is conceptually based on Converse Threads and allows the use of stackless threads called Tasklets (also known as Tasks). In addition, it features complete design flexibility and stackable, independent schedulers. Go [25], developed by Google in 2009, is an object-oriented, concurrent programming language. Generic Lightweight Threads (GLT) was presented in [27], unifying LWT solutions under a unique set of semantics. [An analogous ULT programming model is the task-based oriented solutions such as CompSs \[31\], or a fine-grained OpenMP task employment \[32\].](#)

## 3. Threading Libraries

In this section we describe the two types of threading libraries, OS threads and LWTs, that are analyzed and evaluated in this paper. Moreover, we briefly present the OpenMP PM, for which production implementations are currently based on Pthreads.

For the evaluation of the libraries, from the point of view of OS threads, we have selected Pthreads because it is the standard library that tackles perfectly the current hardware concurrency. In the case of LWTs, Qthreads and MassiveThreads have been selected because these are among the most used lightweight threading models in high-performance computing (HPC); Converse Threads and Argobots were chosen because they correspond to the first (and still currently used) LWT library and the most flexible solution, respectively. Despite Go is not HPC-oriented, we have also included it as representative of the high-level abstracted LWT implementations.

### 3.1. Pthreads API

The Pthreads API [33] defines interfaces and functionality to support multiple control flows—called *threads*—within a process. This API exposes thread management and synchronization primitives to be implemented by threading libraries.

Pthreads offers three PMs that differ in how the threads are bound and which is in control. An important agent in these PMs is the *kernel scheduled entity* (KSE).

KSEs can be managed directly by the OS kernel and the PM changes depending on the threads–KSE mapping.

The *library–thread model* contains a single KSE, and several threads are scheduled and executed on top of it. This relationship is N:1 and may limit concurrency because only a single thread is scheduled at a time. It is leveraged by the GNU Portable Threads library [34].

The *hybrid model* is composed of a set of KSEs, each managing several threads in an M:N relationship. Since LWT libraries follow this hybrid approach, the Pthreads API should be able to accommodate the PM offered by LWTs.

The *kernel–thread model* employs one KSE for each thread that is generated (1:1 relationship). This increases the overhead of the management mechanism because the OS kernel is involved in the scheduling and execution of the threads. This is probably the most used implementation of the Pthreads API that is used in the GNU C library [6].

Pthreads does not expose KSEs as part of the API, although these are present in its execution model. Hence, the Pthreads implementations interpret KSEs differently, leading to the previously discussed mappings between KSEs and threads (N:1, 1:1, or M:N). Therefore, users do not have control over this mapping; instead, they have to follow the mapping offered by the threading implementation. Although some implementations offer functionality for this mapping (e.g. `pthread_setaffinity_np` in the GNU C library), this is not supported by the standard and, therefore, changing the underlying Pthreads implementation may produce a misbehavior in the application/runtime.

### 3.1.1. OpenMP over Pthreads

High-level parallel PMs have been implemented on top of Pthreads in order to promote programming productivity by easing the use of parallel techniques. The most well-known example is OpenMP,

an API that supports multi-platform shared-memory multiprocessing programming. Currently, there exist implementations of OpenMP for most platforms, processor architectures, and operating systems. OpenMP exposes a directive-based PM that helps users to accelerate their codes exploiting hardware parallelism by adding annotations to the code, and at compile time, these annotations are converted to runtime function calls. Intel and GNU develop two commonly used OpenMP implementations that leverage Pthreads in order to exploit concurrency. These runtimes automatically create all the necessary structures and distribute the work among Pthreads.

Since version 3.0, OpenMP supports the concept of tasks, which constitute different pieces of code that may be executed in parallel, and each can be different (e.g., only computation, I/O, communication, etc.). In contrast with work-sharing constructs, distinct OpenMP implementations leverage different mechanisms for task management. In particular, while the GNU version implements a shared task queue for all threads, the Intel implementation incorporates one task queue per thread and integrates an advanced work-stealing procedure for load balancing.

### 3.2. Converse Threads

Converse Threads [12] [13] was one of the first LWT implementations, developed at the University of Illi-

nois in 1996. It is a parallel language-integration solution designed to enable the interaction of different PMs.

Although Converse Threads was designed and developed more than 20 years ago as a general-purpose solution, it is still being used because it composes the underlying layer of the Charm++ implementation [14]. Since its creation, Converse Threads has been extended with several modules (e.g., client–server) that expand the basic functionality and adapt the PM to diverse application scenarios. This continuous development maintains Converse Threads as an appealing solution for HPC environments.

The Converse Threads PM offers two hierarchical thread levels: process (OS threads) and work units. Processes allocate queues where work units are stored. Users may select the number of active processes by means of environment variables.

As an innovative feature, Converse Threads exposes two types of work units: ULTs and Messages. ULTs are the base of the LWT solutions, and represent a migratable (a ULT is executed by an OS thread, paused, and resumed by another OS thread), yieldable, and suspendable work unit with its own stack. A “message” represents a piece of code that is executed atomically. Messages lack an own stack and thus cannot be migrated, yielded, or suspended. Instead, these are recommended for inter-ULT communication, for short nonblocking tasks, and as synchronization mechanisms. Only messages may be inserted into other thread’s queues and this situation reduces flexibility because some codes (e.g., a blocking code) cannot be executed as a message.

Figure 2 depicts the PM offered by Converse Threads, showing the interaction of Converse Threads processes via messages. In that scenario, Process 0 sends a message to Process 1 that is scheduled and executed. Once Process 1 finishes executing the ULTs, it communicates to Process 0 the work completion via the insertion of a message into the Process 0 queue.

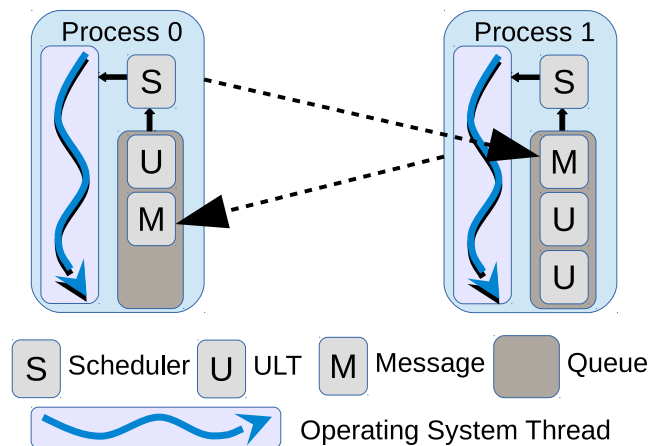


Figure 2: Converse Threads PM and process interaction.

From the point of view of the PM, Converse Threads allows different execution manners, aimed to different scenarios. The behavior is selected with the func-

tion `ConverseInit`, which initializes the environment. If the “normal” mode is selected, threads operate like MPI processes and all the threads execute the overall code. The user is able to select the code portion to execute depending on the thread id. In the “return” mode, `Converse Threads` creates one thread that acts as the master. This thread dispatches the work among other threads by sending messages.

The `Converse Threads` scheduler is a priority system that supports efficiently stackless and standard threads. This scheduler allows two strategies: First-In-First-Out (FIFO) and Last-In-First-Out (LIFO). In order to enhance the flexibility of `Converse Threads`, this library also allows user-defined schedulers that interact with threads.

To complete this concurrent environment, the `Converse Threads` library offers several concurrent implementations of data structures developed specifically for this PM, including queues and lists that are automatically managed by the library.

### 3.3. MassiveThreads


`MassiveThreads` [21] was developed at the University of Tokyo (Japan) in 2014. This LWT library is a recursion-oriented solution that tackles the thread blocking problem when an I/O operation is executed. In addition, this solution provides a tuned load balancing via work-stealing mechanism among threads.

`MassiveThreads` is a consolidated solution in continuous development, which enables it to be used in current hardware systems.


As almost every other LWT solution does, `MassiveThreads` offers two hierarchical levels: Workers (the OS thread) and ULTs. Each worker includes its own work unit queue that is managed by a scheduler. The representation of the PM is illustrated in figures 3a and 3b. The default queue scheduler follows the work-first scheduling policy (Figure 3a): when a new ULT is created, it is immediately executed, and the current ULT is moved into a ready queue. In this scenario, Worker 0 generates a new ULT and the Main task (labeled as *M*) is moved to the queue. Then, *M* may be stolen by idle Workers. Although this policy benefits recursive codes because of the exploitation of data locality, this behavior may be modified to a help-first policy (Figure 3b) at compile time. The help-first policy prevents the worker from executing the new ULTs unless a yield function is called. Therefore, Worker 0 generates a certain number of ULTs that are stored in the queue and Worker 1 steals the lastly created ULT.

The number of workers that are spawned by the `MassiveThreads` environment is selected by the user via the environment variable `MYTH_NUM_WORKERS`. Once the application is launched, this number cannot be modified.

In contrast with `Converse Threads`, `MassiveThreads` does not allow the introduction of work units into other Worker’s queues. Therefore, all the work units are created into the current Worker’s queue and the load balance is pursued with a work-stealing mechanism that allows an idle Worker to gain access to the ready queue of other Workers and to steal a ULT from there. The work-stealing mechanism is also depicted in figures 3a and 3b.



(a) Work-first policy.



(b) Help-first policy.

Figure 3: `MassiveThreads` PMs and scheduler policies.

Once the work units are in the queues, the execution follows the LIFO policy for each worker’s work, and a FIFO policy in case of work-stealing. This algorithm was selected because this scheduling policy is known to be efficient for recursive task parallelism.

`MassiveThreads` includes a mechanism for I/O handling that consists of three procedures, namely, (1) registering a new file descriptor, (2) performing the I/O call, and (3) polling to resume blocked threads. With this mechanism, `MassiveThreads` tackles the blocking thread problem by overlapping communication and computation.

In order to enhance portability from `Pthreads` to `MassiveThreads`, the latter provides a `Pthreads`-like API. This feature allows programmers to convert their legacy codes into `MassiveThreads` applications without effort. Moreover, it allows the use of high-level PMs that are currently written on top of `Pthreads`, with `MassiveThreads` as the underlying library.

### 3.4. Qthreads

`Qthreads` [22] was developed by Sandia National Laboratories (US) in 2008 as a general-purpose LWT imple-

mentation based on the full/empty bit design. The feature that distinguishes this LWT PM is the use of a hierarchy of three levels instead of the two-level structure of the other approaches. The new level is located between the OS thread (called Shepherd) and the work units (ULTs) and it is known as Worker. Shepherds and Workers may be bound to several types of hardware resources (nodes, sockets, cores, or processing units) with the unique restriction that the Shepherd boundary level may lie in a higher level than the Worker level.

Depending on the level of the Shepherds, these may manage one or more Workers. On the one hand, when a Shepherd is bound to a node, it may manage up to  $n$  Workers, where  $n$  is the number of logical cores. On the other hand, when a Shepherd is bound to a logical core, it only manages one Worker bound to the same core. These configurations are determined via a few environment variables; concretely, `QTHREAD_NUM_SHEPHERDS`, `QTHREAD_NUM_WORKERS_PER_SHEPHERD`, `QTHREAD_SHEPHERD_BOUNDARY`, and `QTHREAD_WORKER_UNIT` for the number of Shepherds, number of Workers per Shepherd, and Shepherd and Worker boundaries, respectively. As in the case of `MassiveThreads`, all environment is created inside the initialization function.

Depending on the number of Shepherds (single or multiple) the user is allowed to select the work unit scheduler during the library configuration step. Figure 4 depicts the `Qthreads` system when one Shepherd is bound to a core and one Worker (omitted for simplicity) is spawned per Shepherd. The scheduler configurations integrate work-stealing in order to attain a fair work-load balance among Shepherds. In addition, `Qthreads` enables creating ULTs for specific Shepherds, and those ULTs cannot be stolen by other Shepherds. In Figure 4, Shepherd 1 is not able to steal the last ULT (Assigned ULT), so it steals the previous ULT.

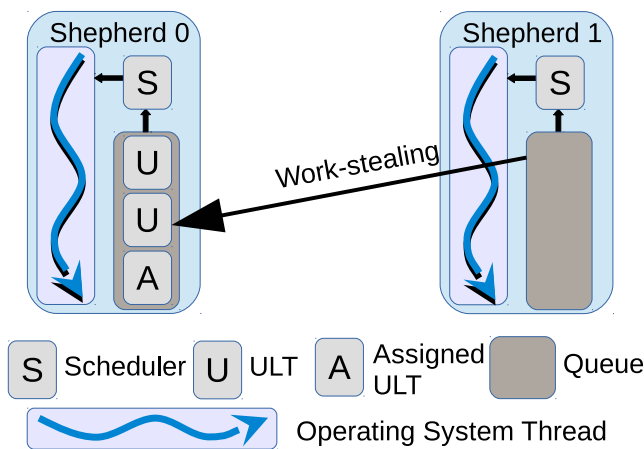


Figure 4: `Qthreads` PM.

`Qthreads` allows that a large number of ULTs access any word in memory. Associated full/empty bits are used not only for synchronization among ULTs but also for leveraging *mutex* mechanisms. This access to memory requires hidden

synchronization, which may severely impact performance.

A large number of distributed structures such as queues, dictionaries, or pools, are offered along with *for loop* and *reduction* functionality. ULT-aware system call functions are also offered in the `Qthreads` API.

### 3.5. Argobots

`Argobots` [7] was developed in 2015 at Argonne National Laboratory (US). It is presented as a mechanism-oriented LWT solution, that means that in addition to be used as a low-level library, it also offers the mechanism for building different environments. Therefore, it allows programmers to create their own PM.

Thanks to its development approach, this PM offers the programmer an absolute control over all the resources. In contrast with previous LWT solutions, the OS threads (named Execution Streams, ES) may be dynamically created at runtime by the user instead of at the initialization point with environment variables. Since those ESs are independent, there is no need for an internal synchronization mechanism. Users may also decide the number of required work unit pools, as well as which ESs have access to each pool. Those pools may be configured with different access patterns depending on the number of producers and consumers. For example, a queue may be accessed by a single ES in order to create ULTs while it may be accessed by several ESs for executing the work units, and vice-versa.




Figure 5: `Argobots` PM using one private pool for ES 0 and a shared pool for ES 1 and ES 2.

Although a default scheduler is defined for each pool, in `Argobots` programmers may create their own instances and apply them individually to the desired pools. The default scheduler implements a LIFO policy and only allowed ESs may interact with the scheduler. Furthermore, `Argobots` supports stackable schedulers, enabling dynamic changes to the scheduling policy that may benefit code portions. The `Argobots` flexibility is represented in Figure 5. This feature enables the programmer to create different environments inside a unique code. As an example, in Figure 5, ES 0 features its own private queue, while ESs 1 and 2 share a work unit queue. This complete flexibility increases the programming difficulty but, at the same time, improves code



The most important features of the threading libraries from the PM perspective are summarized in Table 2.

TABLE 2: Summary of the execution and scheduling functionality offered by the LWT libraries. Pth, Arg, Qth, MTh, CTh and Go identify the threading libraries Pthreads, Argobots, Qthreads, MassiveThreads, Converse Threads, and Go, respectively.

Concept	Pth	Arg	Qth	MTh	CTh	Go
# of Hierarchy Levels	1	2	3	2	2	2
# of work unit Types	1	2	1	1	2	1
Thread Support	✓	✓	✓	✓	✓	✓
Tasklet Support		✓			✓	
Group Control		✓	✓	✓	✓	✓
Yield		✓	✓	✓	✓	
Yield To		✓				
Global work unit Queue	✓	✓				✓
Private work unit Queue	✓	✓	✓	✓	✓	
Plug-in Scheduler	✓	✓	✓	✓	✓	
Stackable Scheduler		✓				
Group Scheduler		✓				

#### 4.1.1. Hierarchical Levels

This number indicates the number of layers inside each PM. Each layer offers its own features and is aimed for different purposes. The main difference between OS threads and LWTs is that the former only contains one level, Pthread itself, while LWTs feature at least two levels. The lowest level in each library is the OS thread representation and in addition it usually includes, in the case of LWTs, a queue for work units and a scheduler. This object receives different names depending on the library: Execution Stream in Argobots, Shepherd in Qthreads, Worker in MassiveThreads, Processor in Converse Threads, and Thread in Go. Users may decide the number of elements to be spawned via environment variables (Group Control row), which is also available to programmers to be set at runtime in the case of Argobots. Pthreads, schedulers, and queues may be created at runtime by the programmer, using the LWT APIs.

Qthreads presents an additional level located in between the two already presented. This level is formed by workers that execute the work units and are managed by the Shepherds. This feature enables further hardware adaptability depending on the combination of hardware and applications.

Work units are at the top level of this hierarchy. These are usually ULTs but, in the case of Argobots and Converse Threads, these may also be a Tasklet. These work units are executed by a OS thread of the lower level or by a Worker in Qthreads.

#### 4.1.2. Work Units

The main work units in threading libraries are ULTs in the case of LWT solutions and Pthreads in the OS threading implementation. Both types of threads are independent, yieldable, migratable codes featuring their own private stack. The main difference is that LWTs are managed in the user space while OS threads are managed by the OS. In addition, Argobots and Converse Threads support

an additional work unit called Tasklet: an atomic piece of code that shares the stack with its executor; in other words, these may be considered as a function pointer. Tasklets (also known as Tasks) are lighter work units than ULTs and are aimed for codes that do not block or context switch.

These work units are stored inside pools/queues waiting to be executed. These structures are also PM dependent and set the library behavior. In the case of Argobots and Pthreads these structures may be created by the user. This increasing difficulty also leads to increased flexibility. Conversely, Qthreads, MassiveThreads, Converse Threads and Go hide that feature to users and the environment may be modified only via environment variables or at compilation time. Go only allows the use of one shared pool while Qthreads, MassiveThreads, and Converse Threads by default assign one queue/pool per OS thread.

#### 4.1.3. ULT Management

Another feature that defines the PM is the functionality that a programmer is allowed to use over the work units. Go is the most restrictive solution. This implementation does not include basic threading operations such as yield, cancel, or resume, that are exposed by the other libraries. Moreover, only one shared queue is employed, and hence the scheduling options are reduced. At the other extreme, Argobots includes not only the common functionality but also an improvement of the yield call. The yield\_to function allows the programmer to pause a current work unit and start (or resume another ULT) without asking the scheduler. In addition, Argobots enables programmers to create their own environment.

All these features make of Argobots the most flexible solution. The other solutions present default schedulers and environment configurations. As a novelty, Qthreads enables programmers to create work units for a specific Shepherd. In the case of MassiveThreads, it only creates work units inside the current Worker’s queue and an internal work-stealing mechanism ensures the load balance. Although some Pthreads implementations allow the use of yield functions, this functionality is not included in the API specification and therefore the corresponding API function names are appended with the \_np suffix that means “non-portable”.

#### 4.2. Usage Difficulty Analysis

Although semantics are important when selecting a library for generating a code solution, the adoption of a PM is also dependent on its ease of use. From this perspective, we classify the presented libraries in three levels [depending on the features that are offered to the programmers for building their application/PM environment](#): easy, medium, and difficult.

In the **easy** level, we include Go and MassiveThreads because of the reduced number of features that depend on the programmer. In the former, the user is responsible of selecting the number of OS threads and then creating and joining the ULTs. In the latter, the programmer also needs to allocate the resources and select if the scheduler may use a Work-first or a Help-first strategy.

In the **medium** level we find Qthreads. In this solution the programmer is involved in more decisions which may



affect performance. Users may select a combination of the number of Shepherds and Workers via environment variables. In addition, users may decide the boundaries of each element (e.g., node, socket, core, or processing unit). Once the environment is set, at the coding level the programmer may consider if the ULTs may or may not be bound to a specific Shepherd.

In the **difficult** level we include the three remaining solutions: Pthreads, Converse Threads, and Argobots, but for different reasons. Pthreads, while not a LWT solution, enforces the use of low-level functionality if the programmer wants to generate an environment where Pthreads share structures and interact among them. Converse Threads offers three distinct models of execution and each model follows its own rules. This feature forces programmers to understand deeply each of the models in order to select the most appropriate approach in each case. For all scenarios, programmers may manage the ULTs in addition to the Tasklets (or messages) for communication among Processors. Argobots offers complete flexibility for environment generation. In addition, this environment may be changed at run time so one part of the application may behave totally different from others. In this LWT library the programmer indicates the number of Execution Streams, the number of pools, the relationship among Execution Streams and pools, which scheduler policy follows each pool, etc. As usual, increasing the flexibility in the library implies more control and more work from the programmers.

## 5. Parallel Code Patterns

Many scientific applications may benefit from the use of OpenMP in order to shorten their execution time. The basic mechanism consists of using OpenMP *pragmas* in order to hint the compiler the code that may be executed concurrently. The compiler translates these directives into OpenMP runtime calls and, at run time, the code is executed in parallel. In this section, we present the most common parallel patterns and explain how current OpenMP runtimes convert the pragmas into parallel code. These patterns will be the basis for our performance evaluation of the different threading solutions.

### 5.1. For Loop

The most frequently used OpenMP directive and probably also the shortest path to produce parallel code is `#pragma omp parallel for`. It may be placed right before a parallel loop composed of independent iterations and produces code where each thread executes its own range of the iteration space. This thread management is transparent to programmers, who are in charge of annotating the parallelizable code with the *pragma*.

Widely-used OpenMP runtimes, `gcc` and `icc`, handle this scenario similarly. The master thread sets a pointer to a function call containing the parallel code in each thread's data structure and it is also responsible for calling the function. All threads wait in a barrier (unless a `nowait` clause is present) at the end of the parallelized loop.

In the case of LWT solutions, the main thread generates one work unit per thread and divides the number of iterations among them. This work unit contains a function pointer to be executed. All the arguments are passed via a structure

```

1 void for_lwt(void * args)
  {
3   arg_for *arg = (arg_for*) args;
   for (int i = arg.ini; i < arg.fini; i++)
5     code(i);
  }
7   ...
   //Main function
   //Allocate memory for structures
9   ULT * lwts[NUM_ULTS]; arg_for * args[NUM_ULTS];

   for (int i = 0; i < NUM_ULTS; i++)
13  {
    //Calculate the number of iterations per LWT
15    ...
    //Arguments initialization
17    args[i].ini = XXX;
    args[i].fini = XXX;
19    //LWT creation
    create_lwt(for_lwt, args[i], lwts[i]);
21  }

23  lwt_yield();

25  //Wait for LWT completion
   for (int i = 0; i < NUM_ULTS; i++)
27    join_lwt(lwts[i]);

```

Listing 1: OpenMP `for` loop parallelism implemented with LWT solutions.

that contains critical information (the number of iterations, variables, etc.) that is necessary to execute the function. Listing 1 shows an abstraction of how the OpenMP `for` loop *pragma* with a static scheduler is translated into LWT code. Lines 1–6 compose the function that is executed inside the ULTs. Lines 8–27 are part of the main function where initially the data structures for ULTs and arguments are allocated (Line 10); Lines 12–21 correspond to the division of the number of iterations among the number of ULTs, the argument initialization, and the ULT creation (line 20). Line 23 allows the main thread to call the scheduler and execute a ready ULT. Once the control is returned to the main function, all the ULTs are joined (line 27) and the work is completed. This example highlights the complexity of the code leveraging low-level LWT APIs.

### 5.2. Task Parallelism

Task parallelism appeared in the OpenMP 3.0 specification as a solution to parallelize unbounded loops and recursive codes. Its usage, however, has been spread to all type of applications that contain pieces of code that may be executed in parallel or that present dependencies among them. In the second case, the runtime generates a directed acyclic graph of tasks and their dependencies and once these are fulfilled, the corresponding task may be executed. OpenMP tasking follows the LWT approach in the sense that tasks are pieces of enqueued code waiting to be executed by an idle thread. This is expressed with the `pragma #pragma omp task`; however, different OpenMP runtimes leverage their own internal approach for tasking. More concretely, `gcc` OpenMP creates a shared task queue for tasks storage. This queue is accessed by all the team's threads. In contrast, `icc` allocates one private queue for each thread in the team. This implementation reduces the contention generated by a shared queue. It also implements a work-stealing mechanism

for load-balancing purposes. This mechanism is used once a thread's task queue is empty and the thread becomes idle. Storing an elevate number of tasks may reduce performance because of contention and the cost of queue reallocation; therefore, `gcc` and `icc` include a non-configurable cutoff mechanism. Once a specific number of tasks is stored (64 times the number of threads for `gcc` and 256 in a thread's queue in the case of `icc`), the new non-dependent tasks avoid the queues and are executed immediately. The situations described in the following two subsections may appear, depending on the code that creates the tasks.

### 5.2.1. Single Region

In this scenario, a thread executing a single or master OpenMP region (`#pragma omp single` or `#pragma omp master`) creates all the tasks in that region. Meanwhile, the other threads in the team are executing them. Once the thread that creates the tasks finishes its work, it joins the others in the task execution process.

The OpenMP implementation in `gcc` uses one shared queue for all threads and the created tasks are pushed into that queue. The threads in the team compete for access to obtain a task. The protection of the queue is enforced via `mutex` and thus contention may increase with the number of threads. Conversely, `icc` OpenMP uses one private queue for each thread. This situation triggers work-stealing because the other threads are idle. The performance in this scenario is also affected by contention because all threads are trying to gain access to the queue. The effectiveness of the work-stealing mechanism may well affect performance.

When using the threading solutions, the main thread generates one work unit per OpenMP task and, as in the `for` loop scenario, the work unit is created with the function pointer and the necessary data.

### 5.2.2. Parallel Region

This scenario occurs when all the threads in a team execute parallel code that creates tasks. During execution, the threads push the new tasks into the task queue (if the cutoff value is not reached), and once this is done, the threads execute the tasks. For `gcc` OpenMP, threads compete to gain access to the shared queue twice: the first time for creating the tasks and the second time for obtaining tasks to execute. In `icc` OpenMP, each thread generates the tasks into its own queue. With this approach, work-stealing is almost nonexistent thanks to a better load balance.

When threading libraries are employed, we have two different parallel levels. The first is mapped to the parallel region, where the main thread generates a work unit for each thread with the function pointer of the region. Then each thread executes the parallel region creating their own work units that are the OpenMP tasks.

## 5.3. Nested Parallel Constructs

When one or more parallel *pragmas* are found inside an already parallel code, it is called nested parallelism. In this case, for the first *pragma*, the runtime spawns a team of threads and for the second *pragma*, each thread in the team becomes the master thread of its new threads' team. In this scenario, the number of active threads grows quadratically. Nested parallelism is not frequent in current applications

because performance drops when the number of threads exceeds the number of cores (oversubscription). However, oversubscription may occur in some situations that the user may be unaware of. For example, a programmer accelerates the code with OpenMP *pragmas* and, in this concurrent code, threads may use external library functions that are parallelized using also OpenMP *pragmas*.

The way OpenMP solutions manage nested parallelism differs. The `icc` OpenMP runtime fulfills the new thread teams reusing idle threads (if any) or creating new threads. Conversely, `gcc` OpenMP does not reuse the idle threads; therefore, each time an OpenMP *pragma* is executed, a new team of threads is created. Since the idle threads are not deleted (but stored in a thread pool), the number of active threads in the system may increase exponentially. To reproduce this scenario, we have implemented code using two nested `for` loops, each with its own `#pragma omp parallel for` directive.

With LWT libraries, the outer `for` loop implementation follows the behavior of that pattern and each work unit executes a range of iterations of the outer loop. Then each work unit creates as many work units as number of threads and divide the iterations of the inner loop among these.

## 6. Microbenchmark Implementation Details

In this section, we discuss how we adapted our microbenchmarks implementing the patterns described in Section 5 to the specifics of each LWT library <sup>1</sup>.

### 6.1. Pthreads

Although `Pthreads` is the supporting library used in production OpenMP and other high-level PMs, we also used this solution as a low-level threading option for our microbenchmarks. In those cases where OpenMP tasks are employed, we generate one `Pthread` for each task and limit the available number of cores by means of the `taskset` command. The thread management relies on the OS itself for all the microbenchmarks.

### 6.2. Converse Threads

In the scenario where `Converse Threads` is used, we employ the "return" mode and leverage `Messages`. The former enables us to follow the OpenMP approach, where there is one master thread while the other processes are treated as slaves. The usage of `Messages` is necessary because these are the only work unit that may be pushed into other threads' queues, and therefore the only way to mimic OpenMP's behavior. In this scenario, the master thread creates as many `Messages` as threads and pushes them to other threads' queues. This limitation, however, prevents the use of `Converse Threads` to support parallel codes because messages cannot yield and hence the requirements of OpenMP are not fulfilled.

### 6.3. MassiveThreads

We have analyzed both `Work-first` and `Help-first` policies. However, only the best of these is shown in Section 7. The difference among these two policies lies on the way a new work unit is treated. While the former pushes the current

1. Available at [https://github.com/adcastel/ULT\\_work/tree/master/lwt\\_microbenchmarks](https://github.com/adcastel/ULT_work/tree/master/lwt_microbenchmarks)

work unit into the queue and executes the recently created work unit, the latter pushes new work units into the queues while the current task continues its execution.

#### 6.4. Qthreads

With its three levels of hierarchy, `Qthreads` accommodates multiple combinations in order to attain high performance in a variety of situations. We have tested a set of combinations, including one Shepherd managing the complete node (it manages up to 72 Workers), one Shepherd per socket (each manages up to 36 Workers), and one Shepherd per core (each manages just one Worker). After a preliminary analysis, we chose two combinations: one Shepherd bound to a node and one Shepherd per core. The first choice is more efficient when there is a reduced number of work units, at the cost of increasing the load imbalance; the second option is more appropriate for scenarios with a higher number of work units. In the presentation of our results we discarded the option with a single Shepherd per socket because it offered lower performance than the other choices for all scenarios.

We also test the functions `qthread_fork` and `qthread_fork_to`, which differ in the work queue where the new work unit is stored. While the former pushes the work unit into the current Shepherd's queue, the latter pushes the work unit into a different Shepherd's queue. If `qthread_fork_to` is chosen, the main thread distributes the work using a round-robin dispatch. Hence, four implementations have been evaluated for each test: `qthread_fork` with one Shepherd per node, `qthread_fork` with one Shepherd per core, `qthread_fork_to` with one Shepherd per node, and `qthread_fork_to` with one Shepherd per core.

#### 6.5. Argobots

The flexibility offered by `Argobots` is two-fold. On the one hand, two different types of work units may be used: ULTs and Tasklets. On the other hand, the work unit pools may be private for each thread or shared among all of them. If the private pool option is selected, the main thread needs to dispatch the work units directly to each thread's pool in a round-robin fashion. Therefore, four possible implementations have been tested. Since Tasklet does not have its own stack and is not yieldable, in those scenarios that require two steps of parallelism (nested and task parallelism), the first of them is performed by using ULTs.

#### 6.6. Go

This library enables only one implementation due to its unique shared work unit queue. All work units need to be pushed into this queue, as the `gcc` OpenMP task implementation does. Therefore, only one possibility is analyzed.

### 7. Performance Evaluation

In this section, we first review the work dispatch/synchronization in both threading solutions and the OpenMP PM. Next, we analyze the different parallel code patterns presented in Section 5 (see implementation details in Section 6). The experiments were performed on an Intel 36-core (72 hardware threads) server composed by two Intel Xeon E5-2695v4 (2.10 GHz) CPUs and 128 GB of memory. GNU's `gcc` 6.1 compiler was used to compile the LWT libraries and OpenMP examples. Intel `icc`

compiler 17.0.1 was used to evaluate the performance of the OpenMP implementations, linked with the OpenMP Intel Runtime 20160808 version. For LWT libraries, we employed `Argobots`, `Converse Threads`, and `Go` libraries updated to 07-2018, `Qthreads` 1.10, and `MassiveThreads` 0.95. All shown results are the average of 500 executions. The maximum relative standard deviation observed in each experiment is between 2% and 5%. Please, note that the microbenchmarks' codes are well-balanced so all OpenMP implementations may offer their best performance.

#### 7.1. Basic Functionality

In OpenMP implementations, a master thread (or work unit) is in charge of creating secondary threads (or work units), and then distributing the work among these slaves. Once that is done, the master completes its work (if any) and waits for the synchronization that may indicate that the overall work is completed. This completion may be enforced using different mechanisms, such as barriers, messages, or thread joins.

Although parallel codes may vary depending on different features, such as granularity, the type of code, or the data locality, the work dispatch and join steps are clearly critical for performance, especially in fine-grained codes.

Figure 7 shows the overhead in terms of time spent creating a single work unit for each thread. In this scenario, the main thread creates and dispatches the work units. As expected, increasing the number of created work units increases the execution time. As an exception, `MassiveThreads` (labeled as MTH) maintains the performance because the new work units are created into the master's own queue. Intel and GNU OpenMP runtimes, labeled as `ICC` and `GCC`, follow the trend of LWT solutions. `Go`'s performance corresponds to the usage of a single shared queue, and therefore, more contention is added when the number of threads is increased.

`Converse Threads` and `Argobots Tasklet`, labeled as `CTH` and `ABT(T)`, employ the tasklets that yield the best performance, thanks to its nature, yielding slightly higher performance than `Argobots` when ULTs (`ABT(U)`) are used and two times faster than the `Qthreads (QTH)` implementation. As expected, results show that creating `Pthreads (PTH)` is more expensive than creating LWTs (excluding the `Go` implementation) because the OS is responsible for managing the creation. The difference between the `Pthreads` and OpenMP implementations appears when all threads are created in a previous parallel section (in OpenMP), and hence the time spent corresponds only to the duration of the work assignment.

Figure 8 displays the time required by the master thread to complete the joining mechanism. Joining mechanisms may differ among implementations. More concretely, `gcc` OpenMP and `Converse Threads` (labeled as `GCC` and `CTH`, respectively) employ a barrier mechanism which makes the time grow with the number of threads. Although `Converse Threads` uses Tasklets, this mechanism does not benefit from it. The marked increase of time in `icc` OpenMP derives from leveraging more than one thread per CPU. The master thread accesses memory allocated by other




Figure 7: Time of creating one work unit for each thread.




Figure 8: Time of joining one work unit for each thread.

threads and therefore the overhead is increased. The remaining libraries use a join mechanism; while Go implements an out-of-order channel communication, Qthreads and Argobots follow a sequential algorithm that checks the status of the memory word or work unit, respectively. The main difference is that Argobots checks the status and frees the work unit structure. MassiveThreads and Pthreads deliver the lowest performance. In the former case, since the main task may be executed by any Worker, each time a thread is joined several checks are triggered. In the latter case, the OS itself waits until the thread has finished and frees the allocated memory.

## 7.2. Code Patterns

In order to maintain a fair comparison among patterns, and at the same time, avoid code modification, we have selected a BLAS-1 function that matches perfectly the fine-grained approach of LWTs and is highly parallelizable. We implement a `scal` function, which multiplies (and overwrites) the components of a vector by a scalar. We avoid coarse-grained codes because then, the thread management overhead is totally hidden by the execution time

In the scenarios where loops (`for` loop and nested `for` loop) are employed, the iterations are divided among the number of threads. In the task parallelism cases, one task is created for each vector element attaining a markedly fine-grained code. This extreme fine level of granularity is chosen in order to understand the behavior of each LWT solution, because this type of parallelism does not hide the thread

management overhead. Conversely, if the execution time of a piece of code is sufficiently long, the overhead is hidden, and therefore there is no significant performance difference between using LWTs or OS threads.

## 7.3. For Loop

For this test, we have created a vector with 1,000 elements that results in a 1,000-iteration `for` loop. Figure 9 illustrates the results. The implementations used in this experiment are Argobots with private pools, Qthreads using a single Shepherd per CPU, and the MassiveThreads Help-first policy, because these attain higher performance than other configurations for each LWT solution. While Argobots (ABT(T) and ABT(U)) presents the highest performance thanks to low creation and join times (see figures 7 and 8), the alternative solutions experiment a notable overhead when increasing the number of threads. Qthreads (QTH) shows a low execution time because of its small join time, but this behavior changes when using more Shepherds than the number of cores (36). Once this number is reached, the total time constantly increases because of the resource sharing overheads. MassiveThreads, Pthreads, and Converse Threads (labeled MTH, PTH, and CTH, respectively) present results 25 times slower than Argobots or `icc`. MassiveThreads suffers from work-stealing and Pthreads from OS management overheads. When Converse Threads uses more threads than physical cores, the performance drops due to synchronization mechanisms. Although Go suffers from contention due to the shared queue, its performance is close to the HPC-oriented solutions thanks to its small overhead in the joining mechanism. The performance of the OpenMP Intel and GNU implementations (labeled as ICC and GCC) is close to that of Argobots and Qthreads, respectively, when all the cores are used.




Figure 9: Execution time of 1,000-iteration `for` loop.

## 7.4. Task Parallelism

For task parallelism, we also used a 1,000-element vector, creating a task per each vector element. Figure 10 exposes the execution time when the 1,000 tasks are created by a single thread in a single region. In this case, the LWT implementations used are Argobots with one private pool per thread, Qthreads using one Shepherd per core, and MassiveThreads with the Work-first scheduler policy. The OpenMP environment has been modified setting the

OMP\_WAIT\_POLICY variable to `passive` to decrease the overhead caused by the contention in the task queue. In this scenario, both `Argobots` work units (Tasklets and ULTs) obtain the highest performance. The reason for this lies on its lighter management mechanisms and its ES independence, which avoid internal synchronization procedures. The elevate number of work units increases the difference between `Argobots` ULTs and Tasklets. This demonstrates that if the code does not need any context switch (e.g. blocking call, communication, I/O), the use of Tasklets benefits performance. `Argobots` Tasklets idea proceeds from `Converse Threads` Messages and hence the proximity of their results. Both stackless units reduce the execution time by a factor of two compared with ULT implementations.

`Converse Threads` attains one of the highest performances (up to 12 times faster than `Pthreads`) thanks to its messages and management that are lighter than the ULT functions. `Qthreads` performs slightly lower (2.5 times slower) than the previous solutions because of two reasons: the use of full-empty bit checks in each memory word and the utilization of more Shepherds than physical cores, which requires additional synchronization. `Go`, `icc`, and `gcc`'s performance are in the middle, and this situation demonstrates that the use of an elevated number of tasks negatively affects those.

The `icc` results demonstrate the effects of the work-stealing mechanism. Workers accessing the master thread's queue to steal work units adds contention. `Go` behaves similarly to `gcc` because both rely on a single shared queue. The lowest performance is attained by `Pthreads` and `MassiveThreads`: the former because we are creating 1,000 OS threads, which causes severe oversubscription; and the latter because the work-first policy implies that, each time a new task is created, the main task is stolen by another thread. As a result, data locality is reduced and there is a drop in the overall performance.




Figure 10: Execution time of 1,000 tasks created into a single region.

Another tasking scenario is where tasks are created inside a parallel region. In that case, each thread creates its own work units. This situation is a two-step algorithm. In the first step, as in the `for` loop scenario, the iterations are divided among the threads; in the second step, the tasks are created.

In this scenario, the choices for each threading solution are the same as in the previous test. Figure 11 displays the

results for this experiment. The use of the two-step algorithm affects negatively `Go` and `Converse Threads`. `Go` suffers the contention added by the shared queue, whereas the synchronization mechanism in `Converse Threads` represents more than 70% of the total execution time. The main reason is that `Converse Threads` needs additional yield calls due to the use of Messages for the first step. `MassiveThreads` is more efficient in this case, because this library is aimed for recursivity. In addition, all the threads in `MassiveThreads` are busy, so the work stealing is almost non-existing. `Qthreads` performance is affected negatively by adding more threads and becomes much slower than other ULT libraries (up to 32 times slower than `Argobots`). Almost all the time difference is due to the join mechanism. Although both `Argobots` implementations use ULTs (that can yield) in the first step, the difference between ULTs and Tasklets is negligible.

On the OS threads side, `icc` offers the highest performance because work stealing has disappeared. This is caused by a perfect load balance. `gcc` outperforms other solutions thanks to its cut-off mechanism (up to eight threads) and to the wait policy value set as in the previous test, attaining results similar to those of `Qthreads`. The lowest performance comes again from the `Pthreads` solution due to the oversubscription caused by the creation of 1,000 threads.

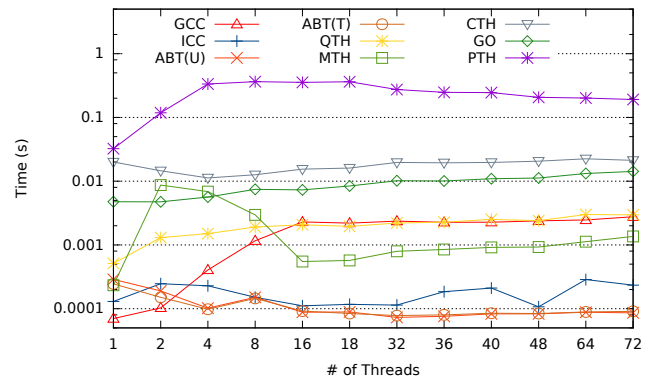


Figure 11: Execution time of 1,000 tasks created into a parallel region.

## 7.5. Nested Parallel Structures

In this case a microbenchmark composed by two nested 1,000-iteration `for` loops has been implemented. The choices for LWT libraries are `Argobots` using a private pool for each thread, `Qthreads` employing one Shepherd per Worker, and `MassiveThreads` with Work-first policy. Figure 12 shows the results for this test. The trend shown by both OS-based approaches, `OpenMP` and `Pthreads`, is different in comparison with that of LWT libraries shown in previous results. This behavior is caused by the sub-optimal implementation of the nested parallel structures in the case of `OpenMP` and oversubscription in the case of `Pthreads`. `gcc` `OpenMP` creates new threads for each nested `pragma` directive and avoids reusing idle threads. As a result, executing this case with 36 threads, `gcc` `OpenMP` spawns 35,036 threads (36 threads for the main team, and 35 threads more for each outer loop iteration). In contrast,




Figure 12: Execution time of a nested parallel for structure with 1,000 iterations per loop.

`icc` OpenMP makes use of the idle threads. However, this approach does not avoid creating a large number of threads. Concretely, it creates 1,296 threads (36 threads for the main team, and 35 for each secondary team). This number is considerably higher than the total number of cores (72), causing severe oversubscription. The `Pthreads` implementation performs close to the `gcc` solution because it follows the same approach. As in previous tests that follow a two step algorithm, `Go` and `Converse Threads` offer low performance. `Go` suffers the contention caused by the employment of just one shared queue. In the case of `Converse Threads`, the addition of `yield` and `barrier` functions slow the execution. However, these still perform higher than `gcc` and `Pthreads` implementations. The three general-purpose solutions (`Argobots Tasklets/ULTs`, `Qthreads`, and `MassiveThreads`) avoid oversubscription by just creating work units instead of OS threads and thus, these yield the highest performance. Avoiding the oversubscription problem reduces the OS thread management overhead, increasing performance with respect to the Intel OpenMP approach by factors of 62, 21, and 25 for `Argobots`, `Qthreads`, and `MassiveThreads`, respectively, with 36 threads.

## 8. Conclusions

We have presented a complete analysis of a set of threading solutions including both OS threads and LWTs. We have performed a PM decomposition of the threading libraries indicating their features. We have proved, by means of experimental tests, that LWTs are aimed for fine-grained parallel codes. For that purpose, we have implemented the most common OpenMP parallel patterns on top of different LWT libraries, and these offer a performance level that is, at least, as good as that attained with `Pthreads` and the OpenMP runtimes. In the case of Intel OpenMP, we have identified some design aspects that may drop performance of common user code patterns. However, current OpenMP solutions were designed for the old concurrent hardware and they are not easy for translating from OS threads to other approaches. These issues may limit its usability in near-future problems.

As a summary, using LWTs in OpenMP critical patterns such as task-parallelism or nested parallel structures may improve the performance compared with the most-used

OpenMP implementations. These scenarios are aimed to tackle the problem of extracting all the computational power of exascale systems.

## Acknowledgments

The researchers from the Universitat Jaume I de Castelló were supported by project TIN2014-53495-R of the MINECO and FEDER, and the Generalitat Valenciana fellowship programme Vali+d 2015. Antonio J. Peña is financed by the European Unions Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 749516. This work was partially supported by the U.S. Dept. of Energy, Office of Science, Office of Advanced Scientific Computing Research (SC-21), under contract DE-AC02-06CH11357.

## References

- [1] B. Nichols, D. Buttlar, and J. Farrell, *Pthreads programming: A POSIX standard for better multiprocessing*. "O'Reilly Media, Inc.", 1996.
- [2] "OpenMP 4.5 specification," [www.openmp.org/](http://www.openmp.org/).
- [3] H. Fu, J. Liao, J. Yang, L. Wang, Z. Song, X. Huang, C. Yang, W. Xue, F. Liu, F. Qiao, W. Zhao, X. Yin, C. Hou, C. Zhang, W. Ge, J. Zhang, Y. Wang, C. Zhou, and G. Yang, "The Sunway TaihuLight supercomputer: System and applications," *Science China Information Sciences*, vol. 59, no. 7, p. 072001, 2016.
- [4] "TOP500 Supercomputer Sites," <http://www.top500.org/>.
- [5] D. Stein and D. Shah, "Implementing lightweight threads," in *USENIX Summer*, 1992.
- [6] "GNU C Library," [www.gnu.org/software/libc/](http://www.gnu.org/software/libc/), Accessed July 2016.
- [7] S. Seo, A. Amer, P. Balaji, C. Bordage, G. Bosilca, A. Brooks, P. Carns, A. Castello, D. Genet, T. Herault, S. Iwasaki, P. Jindal, S. Kale, S. Krishnamoorthy, J. Lifflander, H. Lu, E. Meneses, M. Snir, Y. Sun, K. Taura, and P. Beckman, "Argobots: A lightweight low-level threading and tasking framework," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 3, pp. 512–526, 2017.
- [8] Microsoft MSDN Library, "Fibers," <https://msdn.microsoft.com/en-us/library/ms682661.aspx>.
- [9] "Programming with Solaris Threads," <https://docs.oracle.com/cd/E19455-01/806-5257/6je9h033n/index.html>.
- [10] J. d. Cuvillo, W. Zhu, Z. Hu, and G. R. Gao, "TiNy threads: A thread virtual machine for the Cyclops64 cellular architecture," in *Fifth Workshop on Massively Parallel Processing*, April 2005.
- [11] R. von Behren, J. Condit, F. Zhou, G. C. Necula, and E. Brewer, "Capriccio: Scalable threads for Internet services," in *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles*, October 2003, pp. 268–281.
- [12] L. V. Kale, M. A. Bhandarkar, N. Jagathesan, S. Krishnan, and J. Yelon, "Converse: An interoperable framework for parallel programming," in *Proceedings of the 10th International Parallel Processing Symposium (IPPS)*, April 1996, pp. 212–217.
- [13] L. V. Kale, J. Yelon, and T. Knuff, "Threads for interoperable parallel programming," in *Proceedings of the 9th International Workshop on Languages and Compilers for Parallel Computing*, August 1996, pp. 534–552.
- [14] L. V. Kale and S. Krishnan, *CHARM++: A portable concurrent object oriented system based on C++*. ACM, 1993, vol. 28, no. 10.
- [15] BSC, "Nanos++," <https://pm.bsc.es/projects/nanox/>.
- [16] A. Duran, E. Ayguadé, R. M. Badia, J. Labarta, L. Martinell, X. Martorell, and J. Planas, "OmpSs: A proposal for programming heterogeneous multi-core architectures," *Parallel Processing Letters*, vol. 21, no. 02, pp. 173–193, 2011.
- [17] "GNU Pth - The GNU Portable Threads," <http://www.gnu.org/software/pth/>.
- [18] K. Taura, K. Tabata, and A. Yonezawa, "StackThreads/MP: Integrating futures into calling standards," in *Proceedings of the Seventh ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, 1999, pp. 60–71.
- [19] A. Dunkels, O. Schmidt, T. Voigt, and M. Ali, "Protothreads: Simplifying event-driven programming of memory-constrained embedded

systems,” in *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems*, October 2006, pp. 29–42.

- [20] M. Pérache, H. Jourden, and R. Namyst, “MPC: A unified parallel runtime for clusters of NUMA machines,” in *14th International Euro-Par Conference on Parallel Processing*, 2008, pp. 78–88.
- [21] J. Nakashima and K. Taura, “MassiveThreads: A thread library for high productivity languages,” in *Concurrent Objects and Beyond*. Springer Berlin Heidelberg, 2014, vol. 8665, pp. 222–238.
- [22] K. B. Wheeler, R. C. Murphy, and D. Thain, “Qthreads: An API for programming with millions of lightweight threads,” in *Proceedings of the 2008 Workshop on Multithreaded Architectures and Applications (MTAAP)*, April 2008.
- [23] R. D. Blumofe, C. F. Joerg, B. C. Kuszmaul, C. E. Leiserson, K. H. Randall, and Y. Zhou, “Cilk: An efficient multithreaded runtime system,” *Journal of Parallel and Distributed Computing*, vol. 37, no. 1, pp. 55–69, 1996.
- [24] C. Pheatt, “Intel® threading building blocks,” *Journal of Computing Sciences in Colleges*, vol. 23, no. 4, pp. 298–298, 2008.
- [25] F. Schmager, N. Cameron, and J. Noble, “Gohotdraw: Evaluating the go programming language with design patterns,” in *Evaluation and Usability of Programming Languages and Tools*. ACM, 2010, p. 10.
- [26] “Stackless Python,” <http://www.stackless.com>.
- [27] A. Castelló, S. Seo, R. Mayo, P. Balaji, E. S. Quintana-Ortí, and A. J. Peña, “GLT: A unified API for lightweight thread libraries,” in *IEEE International European Conference on Parallel and Distributed Computing*, Santiago de Compostela, Spain, August 2017.
- [28] A. Castelló, R. Mayo, K. Sala, V. Beltran, P. Balaji, and A. J. Peña, “On the adequacy of lightweight thread approaches for high-level parallel programming models,” *Future Generation Computer Systems*, vol. 84, pp. 22–31, 2018.
- [29] A. Castelló, S. Seo, R. Mayo, P. Balaji, E. S. Quintana-Ortí, and A. J. Peña, “GLTO: On the adequacy of lightweight thread approaches for OpenMP implementations,” in *Proceedings of the International Conference on Parallel Processing*, Bristol, UK, August 2017.
- [30] A. Castelló, A. J. Peña, S. Seo, R. Mayo, P. Balaji, and E. S. Quintana-Ortí, “A review of lightweight thread approaches for high performance computing,” in *Proceedings of the IEEE International Conference on Cluster Computing*, Taipei, Taiwan, September 2016.
- [31] R. M. Badia, J. Conejero, C. Diaz, J. Ejarque, D. Lezzi, F. Lordan, C. Ramon-Cortes, and R. Sirvent, “Comp superscalar, an interoperable programming framework,” *SoftwareX*, vol. 3, pp. 32–36, 2015.
- [32] G. Tagliavini, D. Cesarini, and A. Marongiu, “Unleashing fine-grained parallelism on embedded many-core accelerators with lightweight openmp tasking,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 9, pp. 2150–2163, 2018.
- [33] “Pthreads API,” <http://www.cs.wm.edu/wmpthreads.html>.
- [34] “GNU Portable Threads,” [www.gnu.org/s/pth/](http://www.gnu.org/s/pth/), Accessed July 2016.



**Adrián Castelló** received his BS degree in computer science, the MS degree in advanced computer systems, and his Ph.D. degree in Computer Science from Universitat Jaume I in 2009, 2011 and 2018, respectively. He is a Post-doc researcher at Universitat Politècnica de València and his research interests include deep neural networks, programming models and distributed and shared memory systems.



**Rafael Mayo** received the BS degree from Polytechnic Valencia University in 1991. He obtained his PhD in Computer Science in 2001 at the same University. Since October 2002 he has been an Associate Professor in the department of Computer Science and Engineering in the University Jaume I. Nowadays he is involved in several research efforts on HPC energy-aware systems, cloud computing and HPC system and development tools.



**Sangmin Seo** is a software engineer at Ground X. Prior to that, he worked as an assistant computer scientist at Argonne National Laboratory and a senior engineer at Samsung Research. He received the B.S. degree in computer science and engineering and the Ph.D. degree in electrical engineering and computer science from Seoul National University, respectively. His research interests include high-performance computing, parallel programming models, compilers, runtime systems, and blockchains.



**Pavan Balaji** holds appointments as a Computer Scientist and Group Lead at the Argonne National Laboratory, as an Institute Fellow of the Northwestern-Argonne Institute of Science and Engineering at Northwestern University, and as a Research Fellow of the Computation Institute at the University of Chicago. He leads the Programming Models and Runtime Systems group and his research interests include parallel programming models and runtime systems for communication and I/O on extreme-scale supercomputing systems, and modern system architecture.



**Enrique S. Quintana-Ortí** received the bachelor and Ph.D. degrees in computer sciences from the Universidad Politécnica de Valencia, Spain, in 1992 and 1996, respectively. He is a Professor in Computer Architecture in the Universitat Politècnica de València. Recently, he has participated/participates in EU projects such as TEXT, INTERTWinE, EXA2GREEN and OPRECOMP. His current research interests include parallel programming, linear algebra, energy consumption, transprecision computing and bioinformatics as well as advanced architectures and hardware accelerators.



**Antonio J. Peña** holds a BS+MS degree in Computer Engineering (2006), and MS and PhD degrees in Advanced Computer Systems (2010, 2013), from Universitat Jaume I, Spain. He is currently a Sr. Researcher at Barcelona Supercomputing Center (BSC), Computer Sciences Department. Antonio works within the Programming Models group where he leads the “Accelerators and Communications for HPC” team. Dr. Peña is also the Manager of the BSC/UPC NVIDIA GPU Center of Excellence. He is a Marie Skłodowska-Curie Fellow.