View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by UPCommons. Portal del coneixement obert de la UPC

PARSECSs: Evaluating the Impact of Task Parallelism in the PARSEC Benchmark Suite [1]

Dimitrios Chasapis[*,†], Marc Casas[*], Miquel Moretó[*,†], Raul Vidal[*,†], Eduard Ayguadé[*,†], Jesús Labarta[*,†] and Mateo Valero[*,†]

[*] Barcelona Supercomputing Center (BSC), Barcelona, Spain
[†] Universitat Politècnica de Catalunya – BarcelonaTech (UPC), Barcelona, Spain

In this work we show how parallel applications can be efficiently implemented using task parallelism and we also evaluate the benefits of such parallel paradigm with respect to other approaches. We use the PARSEC benchmark suite as our test bed, which includes applications representative of a wide range of domains from HPC to desktop and server applications. We adopt different parallelization techniques, tailored to the needs of each application, in order to fully exploit the task-based model. Our evaluation shows that task parallelism achieves better performance than thread-based parallelization models, such as Pthreads. Our experimental results show that we can obtain scalability improvements up to 42% on a 16-core system and code size reductions up to 81%. Such reductions are achieved by removing from the source code application specific schedulers or thread pooling systems and transferring these responsibilities to the runtime system software.

Categories and Subject Descriptors: D.1.3 [**Software**]: Concurrent Programming

General Terms: Performance, Measurement, Experimentation

Additional Key Words and Phrases: Parallel Applications, scalable applications, parallel benchmarks, parallel architectures, parallel runtime systems, task-based programming models, concurrency, synchronization

1. INTRODUCTION

In the last few years processor clock frequencies have stagnated, while exploiting Instruction-Level Parallelism (ILP) has already reached the point of diminishing returns. Multi-core designs arose as a solution to overcome some of the technological constraints that uniprocessor chips have, but they exacerbated some others as a counterpart. Multi-core architectures can potentially provide the desired performance by exploiting Thread Level Parallelism (TLP) of large scale parallel workloads on chip. Such large amount of parallelism is managed by the software, which means that the programmer needs to implement highly efficient and architecture-aware parallel codes to achieve the expected performance. This is obviously much harder than programming a uniprocessor chip, which is commonly referred as the *Programmability Wall* [Chapman 2007]. Moreover, dealing with this wall will be even harder in the near future with the arrival of many-core systems with tens or hundreds of heterogeneous cores and accelerators on-chip.

Threading is the most common way to program many-core processors. POSIX threads (Pthreads) [Butenhof 1997] and OpenMP [Chapman et al. 2007] are two of the most common programming models to implement threading schemes. Additionally, MPI [Nagle 2005] can be incorporated to threading codes to handle parallelism in a distributed memory environment. However, to develop efficient threading codes can be a really hard job due to the increasing amount of concurrency handled by many-core processors and the current trend towards more heterogeneity within the chip. Synchronization points are often needed in threading codes to control the data flow and to enforce correctness. However, the cost of these schemes increases with the amount

[1]This paper is published in the journal ACM Transactions on Architecture and Code Optimization (TACO), volume 12, number 4, pp. 41:1-41:22, January 2016. The final publication is available at http://doi.acm.org/10.1145/2829952.

of parallelism handled on chip, seriously hurting performance due to issues like load imbalance or NUMA effects. Also, relaxing synchronization costs often involves significant programming efforts as it requires the deployment of complex and application specific mechanism like thread pools.

Task parallelism [Fatahalian et al. 2006; Blumofe et al. 1995; Bellens et al. 2006; Ayguadé et al. 2009; Tzenakis et al. 2012; Jenista et al. 2011; Planas et al. 2009; Duran et al. 2011] is an alternative parallel paradigm where the load is organized into tasks that can be asynchronously executed. Also, some task-based programming models allow the programmer to specify data or control dependencies between the different tasks, which allows synchronization points relaxation by explicitly specifying the data involved in the operation [Jenista et al. 2011; Ayguadé et al. 2009; Tzenakis et al. 2012; Duran et al. 2011].

The task-based execution model requires to track the dependencies among tasks, which can be explicitly specified by the programmer [Jenista et al. 2011; Zuckerman et al. 2011] or dynamically handled by an underlying runtime system [Duran et al. 2009; Tzenakis et al. 2012; Duran et al. 2011]. When dependencies are detected among tasks, a deterministic execution order is applied by the runtime system to enforce correctness. In this way, all the potential parallelism of the code is exposed to the runtime system, which can exploit it depending on the available hardware. Additional optimizations like load balancing or work stealing [Blumofe et al. 1995; Duran et al. 2011] can be applied at the runtime system layer without requiring any platform-specific consideration from the programmer.

The potential of task-based programming models is expected to be significant in a wide range of areas. The evaluations made so far consider micro-benchmarks [Appeltauer et al. 2009; Podobas and Brorsson 2010], or are limited to specific application domains like graph analysis codes [Adcock et al. 2013] or High Performance Computing (HPC) kernels [Ayguadé et al. 2008; Podobas and Brorsson 2010]. In this work, we aim to evaluate the benefits of task-based parallelism beyond the scope of HPC applications, focusing on a set of parallel applications representative of a wide range of domains from HPC to desktop and server applications. To do so, we apply task-parallelism strategies to the PARSEC benchmark suite [Bienia 2011] and compare them in terms of programmability and performance with respect to the fork-join versions contained in the suite. The main contributions of this paper are:

— We apply task-based parallelization strategies to 10 PARSEC applications.
— We fully evaluate them in terms of performance, considering different scenarios (from 1 to 16 cores) and achieving average improvements of 13%. In some particular cases, the improvements reach 42%.
— We provide detailed programmability metrics based in lines of code, achieving an average reduction of 28% and reaching a maximum of 81%.

The remaining of this document is organized as follows: Section 2 offers background information on the task-based model we chose and the PARSEC benchmark suite. In Section 3 we discuss how these applications are parallelized in Pthreads/OpenMP and then explain our task-based approach. Section 4 shares our experience in programming these applications: the required effort, the versatility of the task-based model, and its current limitations. Section 5 evaluates the performance of the Pthreads/OpenMP codes and our task-based implementations. Section 6 summarizes the related work and, finally, in Section 7 we give our closing remarks.

Table I: PARSEC Benchmark Suite

Benchmark	Description	Native input	LOC
blackscholes	Intel RMS benchmark. It calculates the prices for a portfolio of European options analytically with the Black-Scholes partial differential equation (PDE).	10,000,000 options	404
bodytrack	Computer vision application which tracks a 3D pose of a marker-less human body with multiple cameras through an image sequence.	4 cameras, 261 frames, 4,000 particles, 5 annealing layers	6,968
canneal	Simulated cache-aware annealing to optimize routing cost of a chip design.	2,500,000 elements, 6,000 temperature steps	3,040
dedup	Compresses a data stream with a combination of global compression and local compression in order to achieve high compression ratios.	672 MB data	3,401
facesim	Intel RMS workload which takes a model of a human face and a time sequence of muscle activation and computes a visually realistic animation of the modeled face.	100 frames, 372,126 tetrahedra	34,134
ferret	Content-based similarity search of feature-rich data such as audio, images, video, 3D shapes, etc.	3,500 queries, 59,695 images database, find top 50 images	10,552
fluidanimate	Intel RMS application uses an extension of the Smoothed Particle Hydrodynamics (SPH) method to simulate an incompressible fluid for interactive animation purposes.	500 frames, 500,000 particles	2,348
freqmine	Intel RMS application which employs an array-based version of the FP-growth (Frequent Pattern-growth) method for Frequent Itemset Mining (FIMI).	250,000 HTML documents, minimum support 11,000	2,231
raytrace	Intel RMS workload which renders an animated 3D scene.	200 frames, 1,920×1,080 pixels, 10 million polygons	13,751
streamcluster	Solves the online clustering problem.	200,000 points per block, 5 block	1,769
swaptions	Intel RMS workload which uses the Heath-Jarrow-Morton (HJM) framework to price a portfolio of swaptions.	128 swaptions, 1,000,000 simulations	1,225
vips	VASARI Image Processing System (VIPS), which includes fundamental image processing operations.	18,000×18,000 pixels	127,957
x264	H.264/AVC (Advanced Video Coding) video encoder.	512 frames, 1,920×1,080 pixels	29,329

2. BACKGROUND

2.1. The PARSEC Benchmark Suite

With the prevalence of many-core processors and the increasing relevance of application domains that do not belong to the traditional HPC field, comes the need for programs representative of current and future parallel workloads. The PARSEC [Bienia 2011] features state-of-the art, computationally intensive algorithms and very diverse workloads from different areas of computing. PARSEC is comprised of 13 benchmark programs. The original suite makes use of the Pthreads parallelization model for all these benchmarks, except for `freqmine`, which is only available in OpenMP. The suite includes input sets for native machine execution, which are real input sets. Table I describes the different benchmarks included in the suite along with their respective native input and the lines of code (LOC) of each application. We apply tasking parallelization strategies to 11 out of its 13 applications: `blackscholes`, `bodytrack`, `canneal`, `dedup`, `facesim`, `ferret`, `fluidanimate`, `freqmine`, `streamcluster` and `swaptions` and `x264`. We leave 2 applications out of this study: `raytrace` and `vips`. `Vips` is a domain specific runtime system for image manipulation. Since vips is a runtime itself, it is not reasonable to implement it on top of another runtime system. Therefore we do not include this code in our evaluations. `Raytrace` code has the same extension as `ferret`, `facesim` and `bodytrack` and the same parallel model as `blackscholes` [Cook et al. 2013]. Therefore, since it does not offer any new insight, we do not consider the `Raytrace` code in the paper.

We have a preliminary task-based implementation of the `x264` encoder, which scales up to 14x on a 16-core machine, the same as the Pthreads version. Since we just emulate the same parallel model as the original Pthreads version and obtain the same performance, we do not include this code in the results section as it provides no insight.

2.2. Asynchronous Tasks and Dataflow Model

Tasks offer an easy and abstract way to express parallelism. The OpenMP 4.0 [**?**], a widely used programming standard for shared memory machines, allows the user to annotate functions that can be run asynchronously. It also supports dataflow annotations that describe data dependencies among tasks. This information can be used by the runtime system to synchronize task execution. Standard synchronization schemes are also available (locks, atomics, barriers, etc). In this work we chose to use the OmpSs [Duran et al. 2011] programming model, which is a forerunner of the OpenMP 4.0 tasks. Despite the fact that OmpSs offers advanced features like socket aware scheduling for NUMA architectures or pragma annotations to handle multiple dependence scenarios, in this paper we only use features already available in the OpenMP 4.0 standard. The two models have virtually the same syntax, thus porting OpenMP code to OmpSs and vice versa is straightforward.

```
void load() {
  int i = 0;
  while( load_image(image[i]) ) {
    #pragma omp task in(image[i])
                     out(seg_images[i])
    seg_images[i] = t_seg(image[i]);
    #pragma omp task in(seg_images[i])
                     out(extract_data[i])
    extract_data[i] = t_extract(seg_images[i]);
    #pragma omp task in(extract_data[i])
                     out(vectoriz_data[i])
    vectoriz_data[i] = t_vec(extract_data[i]);
    #pragma omp task in(vectoriz_data[i])
                     out(rank_results[i])
    rank_results[i] = t_rank(vectoriz_data[i]);
    #pragma omp task in(rank_data[i])
                     out(outstream)
    t_out(rank_data[i], outstream);
    i++;
  }
  #pragma omp taskwait
}%
```

Fig. 1: `Ferret` implementation in OmpSs

Figure 1 shows a simplified version of the `ferret` benchmark implemented in OmpSs, an application that is parallelized with a pipeline model. The programmer can use pragma directives to identify functions that should run asynchronously. These task pragmas can have dataflow relations expressed with the use of `in`, `out` and `inout` annotations. These declare whether a variable is going to be read, written or both by the task. An underlying runtime system is responsible for scheduling tasks, track dependencies, balance the load among available threads and ensure correct order of execution, as dictated by the dataflow relations. In our example data dependencies will force tasks spawned in the same iteration to run in sequential order, while tasks from

different iterations can run concurrently. An exception is `t_out` which shares a common output between all instances, `outstream`, to store the final results of `ferret`.

3. APPLICATION PARALLELIZATION

In this section we discuss how the PARSEC applications are parallelized in Pthreads/OpenMP2.0 and how they can be implemented efficiently using a task-based approach. When possible, we exploit dataflow relations in order to take advantage of implicit synchronization (as described in Section 2.2). If it is not possible we use conventional synchronization primitives such as locks, atomics and barriers.

Blackscholes. This application solves a Black-Scholes Partial Differential Equation [Black and Scholes 1973] to calculate the prices for a portfolio of ten million European options.

Pthreads. This version simply divides the portfolio into work units by the number of available threads, and stores them into the `numOptions` array. Each thread calculates the prices for its corresponding options and waits in a barrier until all the threads have finished executing. The algorithm is run multiple times to obtain the final estimation of the portfolio.

Task-based. In the case of the task-based version, we divide the work into units of a predefined block size. This block size allows having much more task instances than threads, which implies a much better load balance, as this is an embarrassingly parallel application with no dependencies among tasks in the same run.

Bodytrack. Computer vision application that tracks a marker-less human body using multiple cameras through an image sequence. The application employs an annealed particle filter to track the body using edges and the foreground silhouette as features of interest.

Pthreads. `Bodytrack` applies the same algorithm on each frame of the image sequence to track the different poses of the body. The human body is modeled as a tree-based structure, consisting of 7 conic cylinders. It reads 4 images taken from several cameras to capture a scene from 4 different angles, thus each frame consists of these 4 images. These images are read and encoded to a single data structure. For each frame, `bodytrack` extracts the edges and silhouette features for each of these 4 images. In this feature extraction stage we have 3 different kernels.

(1) **Edge detection**: Gradient based edge detection.
(2) **Edge smoothing (phase 1)**: Gaussian filter used to smooth edges applied on array rows.
(3) **Edge smoothing (phase 2)**: Gaussian filter used to smooth edges applied on array columns.

Afterwards, `bodytrack` goes through an annealed particle filter stage, which consists of M annealing layers over a set of N particles. The particles are multi-variate configurations of the state and location of the tracked body. Given the image features, the particles are assigned weights, which increase or decrease the chance that a particle represents a body part. N particles are then chosen, depending on the probability dictated by their weights. Random noise is added to this set of particles, creating a new set. This process is repeated for all annealing layers. `Bodytrack` then picks one of the M configurations, the one which has the highest weighted average. This process has two parallel kernels.

(4) **Calculate particle weights**: Computes weights for the particles, using the edges and silhouette produced from the previous stages.

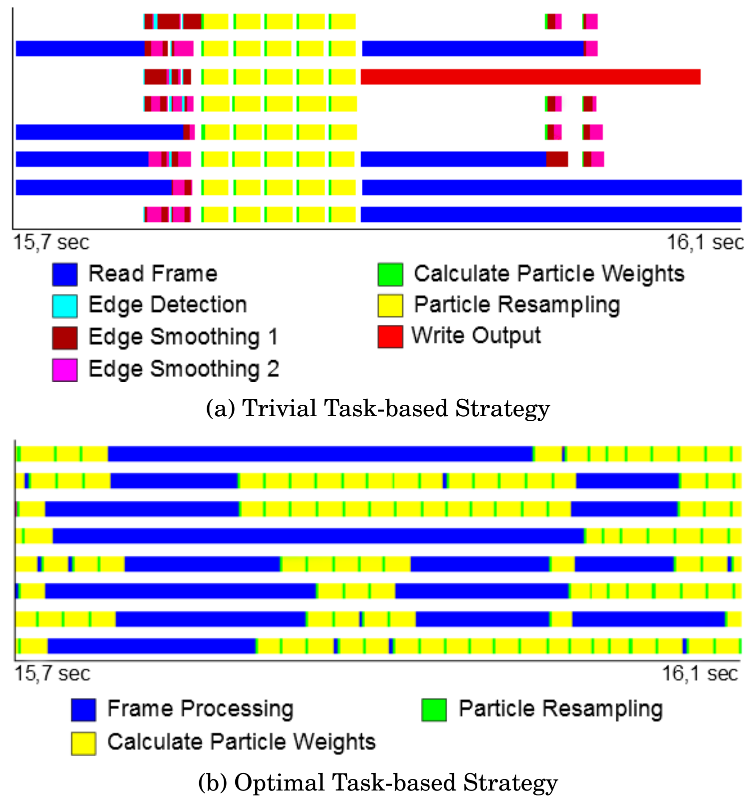

(a) Trivial Task-based Strategy

(b) Optimal Task-based Strategy

Fig. 2: Parallel execution of Pthreads and task-based versions of `bodytrack` on an 8-core machine and native input size. Different parallel regions correspond to different colors. White gaps in the figure, represent idle time.

(5) **Particle Resampling**: Adds Gaussian random noise to the particles, thus creating a new set of particles.

In the case of Pthreads, the 4 images of a frame are read and processed in parallel using one thread per image. The Pthreads implantation is limited by the 4 images it can process concurrently, while there is no other candidate work at this point. A specific asynchronous I/O implementation is required to read the files in parallel. Then, the features extraction stage is executed using all the available threads, with a synchronization barrier at the end of each phase. The same structure is followed in the annealed particle filter stage, with two barriers at the end of each phase. Between the two stages, serial code has to be executed, which leaves only one thread busy and the rest idle. Finally, the output results are written sequentially in one file.

Task-based. In the case of the task-based version, we adopt a more coarse grain approach. We do not parallelize the feature extraction stage, instead we taskify the whole frame processing, allowing concurrent execution of all frames. The parallel kernels of the annealed particle filter stage are taskified in our version, and synchronization is achieved by dataflow annotations. Each frame needs to be written when calculations are completed. In our version we can do this asynchronously while the threads are busy

with the processing stage of another frame. Thus, output I/O is effectively overlapped with computation stages.

Figure 2 shows parallel executions of two different task-based implementations: The first one just mimics the Pthreads behavior (2a) and the second is an optimal task-based implementation (2b). Different colored boxes represent different task types, as well the duration of that task type on each core. In both cases, the white gaps denote the time each thread spends idle. Both figures show the same duration for each execution. In the optimal version, all functionality is implemented within the frame-processing task, thus execution time for read-frame, edge-detection and edge-smoothing is represented with blue color (frame-processing). Tasks particle-resampling and calculate-particle-weights are also implemented as nested tasks. They are displayed with different colors (green and yellow respectively). We can observe that the Pthreads-like version suffers from greater idle time compared to its optimal task-based counterpart. Work is distributed more efficiently in the optimal implementation by processing different frames concurrently. This allows us to overlap I/O and serial code segments of one with available work from another one.

Canneal. This kernel uses a cache-aware simulated annealing [Banerjee 1994] to optimize routing cost of a chip design. `Canneal` progressively swaps elements that need to be placed in a large space, eventually converging to an optimal solution. The problem is stored as a list with routing costs between nodes.

Pthreads. This version compares random element pairs of the graph concurrently and swaps them until it converges to an optimal solution. No locks are used to protect the list from concurrent accesses/writes, but swaps are done atomically instead. However, the evaluation of the elements to be swapped is not atomic. This means that disadvantageous swaps may occur, which will require the algorithm to eventually swap them again. This method has provided better results than the alternative algorithm with locks [Bienia et al. 2008].

Task-based. Our task-based version follows the same paradigm. Several tasks are spawned without any dependencies between themselves. We use the same atomics as with the Pthreads version. Since tasks work with an arbitrary number of list elements, it is not possible to describe which elements of the list a task is going to randomly access.

We also try an alternative fine grain implementation, where a task is spawned for each random pair of list elements. This would allow the runtime to know if two tasks are working on the same list of elements. However this implementation implied fine-grain tasks. Each task would merely do a single swap between two list elements. The overhead of the dynamic scheduling is a problem in this scenario. A more complex but more efficient solution is suggested by Symeonidou et al. [Symeonidou et al. 2013] with the use of memory regions. Adopting this method in a task-based model would allow the programmer to describe parts of the list (or other pointer based data structures) and express dataflow relations as abstract memory regions. This solution also implies fine-grain tasking and is not evaluated at the Symeonidou's work.

Dedup. The `dedup` kernel is used to compress data streams using local and global compression to achieve higher compression rates. This method is called deduplication [Quinlan and Dorward 2002].

Pthreads:. Dedup is parallelized using a pipeline model with the following stages:

— **Fragment**: First, the data-stream is read and partitioned at fixed positions into coarse grain data chunks. Each chunk can be processed individually by the rest of the stages. This stage is executed on a single thread.

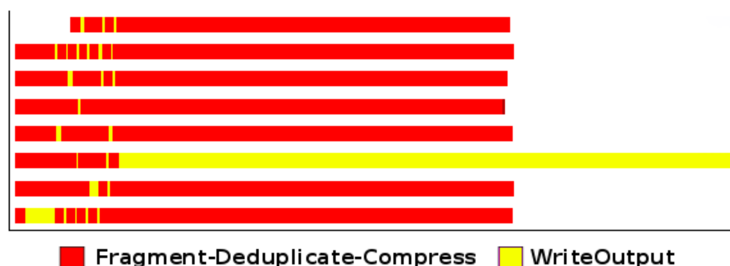

Fig. 3: Parallel execution of the task-based version of `dedup` on an 8-core machine and native input size. Different task types correspond to different colors.

— **Fragment Refine**: A new data chunk initiates the second pipeline stage, where it is further partitioned into smaller fine-grain chunks. The portioning is done by using the Rolling-fingerprint algorithm.
— **Deduplication**: This stage eliminates duplicate fine-grain chunks. Unique chunks are stored in a hash-table. Locks are used here to protect each bucket from concurrent accesses.
— **Compress**: At this stage chunks are compressed in parallel. Identical chunks are compressed only once as duplicates are removed at the deduplication stage.
— **Reorder**: This stage writes the final compressed output data to a file. It writes only unique chunks' compressed data and for the duplicates it stores their hash values. However this stage needs to reorder the data chunks as they are received to match the original order of the uncompressed data.

The Pthreads version maintains a queue and a thread pool dedicated to each stage. When a chunk becomes available at one stage, it is moved to the queue of the next stage. Each stage polls at its queue for available chunks to process. The reorder stage is done sequentially with a devoted thread that can be in an idle loop waiting for previous stages to finish. Each thread pool comprises by a number of threads equal to the number of available cores. The only exceptions are `Fragment` and `Reorder` stages, which are served by a single thread each.

Task-based. In our implementation we taskify each pipeline stage and express data dependencies using static arrays and dataflow relations, one for each pipeline stage. `FragmentRefine` however partitions the data chunks into very fine grain segments, ranging from a few hundreds to thousands. For such granularity, our approach suffers from high overheads due to dynamic schedulnig overhead. The same is observed in [Vandierendonck et al. 2013], where an alternative approach is adopted. In their approach, two pipelines are identified: The outer pipeline, consisting of stages `Fragment`, `InnerPipeline` and `Reorder`. The inner pipeline consists of `FragmentRefine`, `Deduplicate` and `Compress`. To reduce the dynamic scheduling overhead, they merge together `Deduplicate` and `Compress`. By doing so, the available parallelism is limited, but still there is enough work not to harm performance and scalability. In our approach, we merge together the inner pipeline, creating one sequential function, exploiting only the parallelism available in the outer pipeline. Even in this scenario, the available parallelism is still abundant, since the application is bound by the writing of the output file, which is sequential. Figure 3 shows a trace of the task-based version. We can see that communication stage (in yellow) is effectively overlapped with the computation stage (in red), however, there is not enough work to keep all the threads finish, until the end of the execution.

Furthermore, we modify the Reorder stage, by replacing it with a simple stage where the chunk is simply written to file (WriteOutput). Using dataflow relations and a shared output resource between the WriteOutput tasks, we ensure that chunk N-1 will be written before chunk N. Thus, we do not need to reorder data chunks in this task type. Moreover, the scheduler makes sure chunks are written as soon as they become available by the InnerPipeline task, an improvement over the Pthreads version, where Reorder instances need to idle wait until all previous chunks ones have been written. Another difference between the two versions, is that Pthreads oversubscribe threads to cores for each pipeline stage, while in our implementation we only assign one thread to each core.

Facesim. Computes a visually realistic human face animation by simulating the underlying physics. As input it uses a 3D model of a human face containing both a tetrahedra mesh and triangulated surfaces for the flesh and bones, respectively. Additionally it uses a time sequence of muscle movement [Sifakis et al. 2005].

Pthreads. The application statically decomposes the original tetrahedron mesh into smaller partitions, equal to the number of available threads. There are three main parallel kernels:

— **Update State**: Calculates the steady properties of the mesh, constrains like stress and stiffness.
— **Add Forces**: Computes the force contribution between vertices acting on the 3D model.
— **Conjugate Gradient (CG)**: An iterative method that solves the linear system produced by the other two previous kernels and find the final displacement of the vertices for the current frame.

Update_State and Add_Forces kernels consist of one and two parallel loops respectively, while CG has three. Synchronization between loops and kernels is achieved by barriers. The corresponding force computations from the skeleton are also done in Update_State and Add_Forces, but after the parallel computations on the tetrahedra mesh have been made. In Pthreads a master thread is assigning work to all threads in a round-robin fashion through a queuing system. Each thread maintains its private queue, which is protected by locks.

Task-based. In the task-based version, the application level queuing system is completely replaced by the OmpSs runtime. In our initial implementation all parallel loops are taskified. Additionally, in Update_State there is a sequential code segment, Update_Collision_Penalty_Forces. This code segment operates on the bones, while the parallel loop of Update State does so on the tetrahedra. By taskifying it and adding dataflow relations between this section and the following Add_Forces kernel, we can overlap Update_Collision_Penalty_Forces with the rest of Update_State.

To improve performance we refactor tasks' creation in CG by nesting the first task creation loop inside another task. This enables us to overlap task creation time with computation, which contributes to increase Facesim's task-based implementation performance. Although we achieve better scalability than the original code, task creation still imposed overheads. To address this issue we replace tasks in CG with the OmpSs parallel loops construct (equivalent to the OpenMP one), which implements loop worksharing with a task. Even though this approach limits the available parallelism (barrier synchronization, no dataflow annotations), the overhead associated to task creation and scheduling is greatly reduced and overall performance improved.

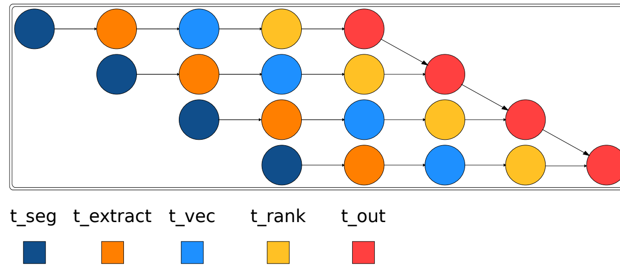

Fig. 4: Task-graph of ferret showing the pipelined execution model. Edges show data dependencies among different tasks.

Ferret. Content similarity search server for feature-rich data [Lv et al. 2006] like audio, video, images, etc. The benchmark application is configured for image similarity search.

Pthreads. Ferret is parallelized using a pipeline model. A serial query is broken down into 6 pipeline stages:

— **Load**: This stage loads an image that is going to be used as a query.
— **Segmentation**: At this stage the image is decomposed into the different objects displayed on it. Different weight vectors are to be assigned on each object to achieve better results.
— **Extract**: At this stage a 14-dimensional vector is computed for each object from the segmentation stage, describing features such as color, area, state, etc.
— **Vectorization**: This is the indexing stage that tries to find a set of candidate images in the database.
— **Rank**: This stage ranks the results found, using the EMD metric for each query-object's vector and the database's image vectors.
— **Output**: Outputs the result of the ranking stage. Multiple instances of this stage need to run serially, since they all share the same output stream.

In the Pthreads version every stage is served by a dedicated thread-pool of N threads each, where N is the number of available cores. The only exceptions are the `Load` and `Output` stages that are executed by a respective single thread. Each stage polls on its corresponding queue for available work. When a stage finishes, it pushes the results to the next stage's queue.

Task-based. In this version, we implement a variation of this pipeline model. As soon as the first stage, `Load`, finds a new image, it spawns all stages of a pipeline for that image, thus reducing the pipeline to five stages. We model the dataflow relations between different stages as simple one dimension arrays, as shown in Figure 1. Tasks working on different image queries do not share any dependencies. An exception is task `t_out` which shares the same output file between all pipelines, thus sequential execution is forced between all instances of this task. The pipeline stages and dependencies are constructed a priori, which is good enough for this application, but this is not always the case. [Lee et al. 2013] proposes a system that can handle dynamic pipeline creation by constructing a DAG with the stages using indexes and the `cilk_continue` and `cilk_wait` keywords. Indexes are used to define the different pipeline stages, while `cilk_continue` creates a stage that can run once all previous stages in the same pipeline iteration are done, and `cilk_wait` creates a stage that will wait for its stage counterpart of the previous iteration to finish. A strategy based on versioning the dependency objects between the stages has been proposed [Vandieren-

donck et al. 2011]. Output dependencies are renamed and privatized, thus the static array for privatization is not required.

Figure 4 shows the task-graph of the `ferret` application. Colored nodes denote the concurrent tasks (each color matches a specific task type). Tasks that have data dependencies are connected by directed edges. By inspecting the task-graph we can see a pipeline pattern of execution. Despite the fact that the task-based approach does not significantly improve the overall performance, as we can see in Section 5, it significantly reduces the effort required to express the pipeline parallelism, compared its Pthreads counterpart, as it is shown in section 4.1 in detail.

Fluidanimate. This application simulates incompressible fluid interactive animation, using the Smoothed Particle Hydrodynamics (SPH) method [Müller et al. 2003].

Pthreads. `Fluidanimate` uses five special kernels which are responsible for rebuilding the spatial index, computing fluid densities and forces at given points, handling fluid collisions with the scene geometry and finally updating particle locations The fluid surface is partitioned and each thread works on its own grid segment. The kernels are parallelized as do-all loops, separated by barriers. Moreover, there are cases where these threads need to update values beyond their partition, which are handled using locks.

Task-based. The task-based implementation follows the same approach, we apply a loop tiling transformation, for each parallel loop, and taskified each iteration. We maintain the same barrier and lock synchronization scheme, using the OmpSs synchronization primitives.

Freqmine. Data mining application that makes use of an array-based version of the Frequent Pattern (FP) growth method for Frequent Itemset Mining [Grahne and Zhu 2003].

Pthreads. The application uses a compact tree data structure, denoted *FP-tree* [Han et al. 2000], to store information about frequent patterns of the transaction database. The FP-tree is coupled with a header table, which is a list of database items, sorted by decreasing order of occurrences. The FP-growth algorithm traverses the FP-tree structure recursively, constructing new FP-trees until the complete set of frequent itemsets is generated. There are three parallel kernels. The `Build_FP-tree_header_table` kernel performs a database scan and counts the number of occurrences of each item. The result is the FP-tree header table. `Build_Prefix_tree` kernel performs a second database scan required to build the prefix tree and the `Data_Mining` kernel obtains the frequent itemset information by using the previous two structures. It creates an additional lookup table, which allows faster traversals on sparse itemsets. The original PARSEC benchmark uses OpenMP2.0 for loop parallelization inside each kernel.

Task-based. In our implementation we taskify each iteration. We do not use any dataflow relations in this application, and resolve to adopt the locking and barrier synchronization used in the original OpenMP version.

Streamcluster. Streamcluster is a kernel that solves the online clustering problem. It takes a stream of points and then groups them in a predetermined number of clusters with their respective centers.

Pthreads. Up to 90% of total execution time is spent in function `pgain`, computing whether opening a new center is advantageous or not. For every new point, function `pgain` calculates the cost of making it a new center by reassigning some points to it and comparing it to the minimum distance $d(x, y) = |x - y|^2$ between all points x and y. The result is accepted if found to favor the new center. Data points are statically partitioned

by a given block size, which determines the level of parallelism in the application. In the Pthreads version this is equal to the number of threads.

Task-based. In our implementation we follow a different decomposition strategy, making the number of tasks independent of the number of partitions. Barriers are employed to synchronize accesses to a partition in both Pthreads and the task-based implementation. In the case of Pthreads, an additional user implemented library is used for the barriers. This library is not required in the case of the OmpSs implementation, as the runtime already has a generic barrier implementation.

Swaptions. Economics application that uses the Heath-Jarrow-Morton (HJM)[Heath et al. 1992] for pricing of a portfolio of swaptions. To calculate prices it employs the Monte Carlo simulation.

Pthreads. The application stores the portfolio into an array. In the Pthreads version, this array is divided by the number of available threads, each thread working on its own part of the array.

Task-based. We use the exact same strategy, where each task works on a part of the array. No data dependencies exist between the tasks.

4. PROGRAMMABILITY

Different models and languages offer diverse ways to express concepts, such as parallelism or asynchrony. In this section we evaluate how successful and easy it is to express parallelism using task-based models. A good proxy to evaluate how complex a particular implementation is the number of lines of code it takes. Despite being a metric proposed some decades ago, comparing different programming models in terms of the total number of code lines is still a valid metric. Indeed, recent publications make an extensive use of it [Vandierendonck et al. 2011; Dongarra et al. 2008].

4.1. Lines Of Code

The reduction of the lines of code (LOC) attests to a more compact and readable code. In some of our PARSEC task-based implementations, a simple pragma directive replaced application specific schedulers, scheduling queues, thread pooling mechanisms and lock synchronization. We do not change the algorithm in any of the task-based parallel strategy implemented in the PARSEC suite. Figure 5a shows a normalized comparison between the lines of code of our task-based implementations and the original Pthreads/OpenMP implementations of the PARSEC 3.0 distribution. The PARSEC 3.0 versions we refer to are always the Pthreads versions, except in the case of `freqmine` where, since there is no Pthread version available, the OpenMP2.0 version is taken as reference. We preprocess all source files so that they only contain lines of code relevant to the respective programming model[2]. Figure 5b shows the total lines of code comparison when we only consider files that are relevant to the parallel implementation, that is, files that contain calls to Pthreads or task invocations, asynchronous I/O implementations, atomic primitives, etc. In this graph we see that the reductions in terms of lines of code of our task-based strategies are significant. In case of `bodytrack`, we are able to remove 81% of the code lines. Since `Bodytrack` implements its own scheduler to deal with load balancing, there is much room for code reductions by replacing this ad-hoc mechanisms for a few pragma annotations.

By using tasks and dataflow relations, it is very easy to implement pipelines. We adopt this approach for both `dedup` and `ferret`, which result in a significant decrease

[2]PARSEC benchmarks contain mixed serial, Pthreads, OpenMP and TBB source code, and make use of macros to enable conditional compilation for only one programming model at a time.

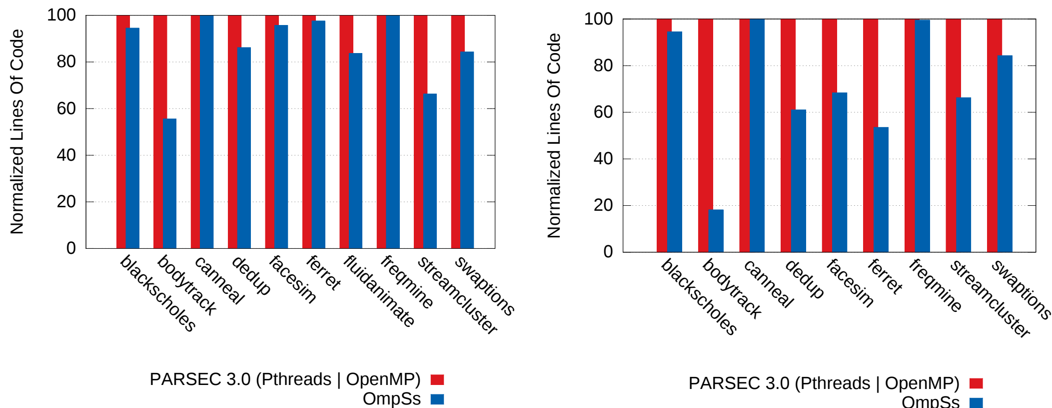

(a) Comparison between all source files.

(b) Comparison between only source files containing parallel code.

Fig. 5: Comparison of lines of code between our task-based implementations and the original Pthreads or OpenMP versions.

in LOC (38% and 46%, respectively). Figure 1 shows the pipeline code for `ferret`. All that is required is to taskify the different pipeline stages and make sure that dataflow relations force in-order execution of tasks in the same pipeline instance. The Pthreads version requires the implementation of queues between each stage, which must also be safe to use by multiple threads and concurrent accesses. `Streamcluster` and `fluidanimate` task-based versions also reduce lines of code by 33% and 21% respectively, by removing the need for an additional, user implemented, barrier library. `Blackscholes` and `swaptions` are relatively simple applications, containing only one do-all parallel loop each. In these cases the LOC difference is minimal (0.5% and 15%, respectively). In the cases of `canneal` and `freqmine` we see no difference in LOC. `Canneal` is not a data parallel application and in both cases Pthreads and tasks are used merely as thread launching mechanism, while the synchronization effort is essentially the same.

It is worth noting that conventional synchronization primitives can still be used with tasks, without penalizing the programmer. `Freqmine` is implemented in OpenMP, which excels at parallelizing loops with very little effort from the programmer and is the ideal programming model for this application. In our implementation we simply taskify the loops, essentially not affecting LOC. `Facesim` also benefits from the tasks-based approach by 37%, as the queues required to schedule work have been completely removed. Overall, we see that the task-based model reduces code size and by 28% on average.

5. PERFORMANCE EVALUATION

In this section we compare our task-based implementations to the original PARSEC implementations in Pthreads or OpenMP.

5.1. Experimental Setup

The experiments are performed on an IBM System X server iDataPlex dx360 M4, composed of two 8-core Intel Sandy Bridge processors E5-2.60Hz, 20MB of shared last-level cache. There are eight 4GB DDR3 DIMM's running at 1.6GHz (a total of 32GB per

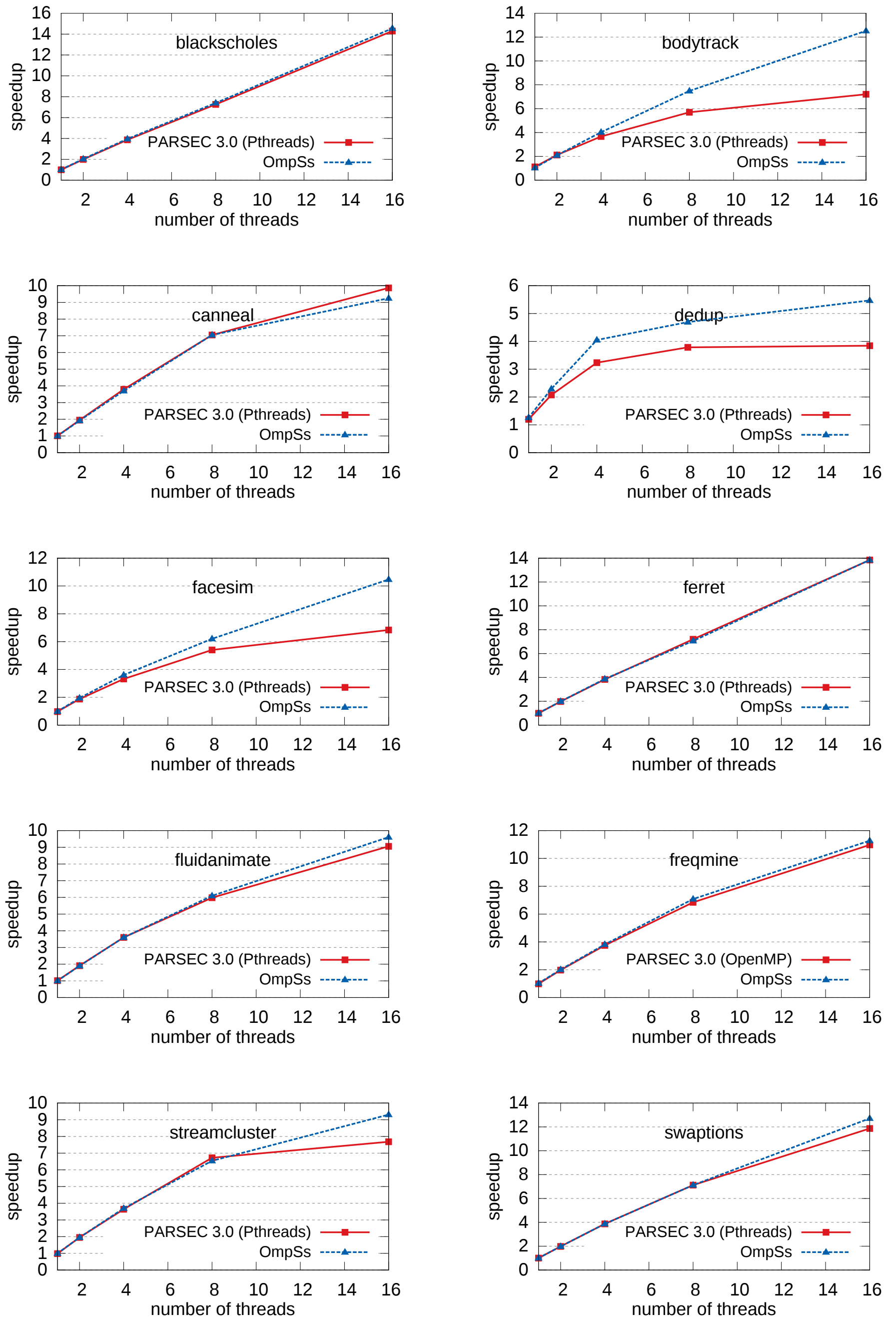

Fig. 6: Comparison of scalability between the task-based implementations and the original (Pthreads/OpenMP) versions.

node and 2GB per core). The hard drive is an IBM 500GB 7.2K 6Gbps NL SATA 3.5. We make use of the OmpSs programming model [Duran et al. 2011] and its associated toolflow: nanos++ runtime system (version 0.8a), Mercurium source-to-source compiler (version 1.99), and gcc 4.7 as the back-end compiler. We run all benchmarks using their respective native inputs as described in Table I.

5.2. Scalability

Figure 6 shows how the task-based codes scale compared to the PARSEC Pthread-/OpenMP versions. Results are shown individually per benchmark as we increase the number of cores assigned to the application and normalized to the execution time of the serial implementation [3] of the application. Nearly all applications scale linearly up to 4 cores.

In the case of `bodytrack`, as described in Section 3, by concurrently executing different frames, there is always enough work for all threads, while by taskifying the output stage of each frame, we overlap this I/O bottleneck with other computation stages. The speedup when run on 16 cores is 12.1x, while the Pthreads implementation reaches a poor 6.8x speedup when run on 16 cores. The `dedup` application has a very expensive stage that writes the compressed data to the output file. Our task-based implementation is very effective in overlapping this time with computation from the compression stage. Also, the task-based version does not have to reorder the data chunks, since the I/O execution takes place in-order as dictated by dataflow relations. This results in an impressive 30% performance improvement of the OmpSs version with respect to Pthreads when run on 16 cores. The Pthreads `facesim` implementation is burdened by barriers that limit available parallelism. By using dataflow relations we taskify sequential segments of significant cost we effectively synchronize them with parallel sections preceding and following it. The performance improvements comes from the overlap of sequential computations with parallel sections. The task-based parallelization of `facesim` reaches a speedup of 10.2x when run on 16 cores, while the PARSEC code only reaches a 6.4x speedup.

In the cases of `blackscholes`, `canneal`, `ferret`, `fluidanimate`, `freqmine` and `swaptions`, the Pthreads/OpenMP versions already achieve good scalability results. With the exception of `ferret`, the task-based codes have very close resemblance to their Pthreads/OpenMP counterparts, and have offered reduced opportunities for OmpSs to dynamically exploit additional parallelism. The parallel implementation in these applications, with the exception of `ferret`, is limited to parallel do-all loops with barrier synchronization, essentially exploiting the same amount of parallelism among all versions (OmpSs/Pthreads/OpenMP).

In the case of `ferret`, although the code is substantially different, both versions employ the same pipeline model and deliver the same level of parallelism, which is already high in the Pthread version. We express a bit of extra parallelism by extending the pipeline with multiple stages, which write to the output file, effectively overlapping some communication with computation. However, the final impact in the total execution time is limited as the time needed to write the output file is a very small fraction of the total execution time. Finally, we observe performance gain (18%) in `streamcluster`, which can be partly attributed to the more efficient barrier implementation of OmpSs, when compared to the user implemented barriers of the Pthreads version. However, the most important performance drawback that the original Pthreads implementation suffers from, is the negative NUMA effects. This issue is observed when we run our

[3]The PARSEC benchmark suite provides a serial implementation for `blackscholes`, `bodytrack`, `dedup`, `ferret`, `freqmine` and `swaptions`. For the other benchmarks, the original Pthreads parallel implementation executed on a single core is considered as the baseline.

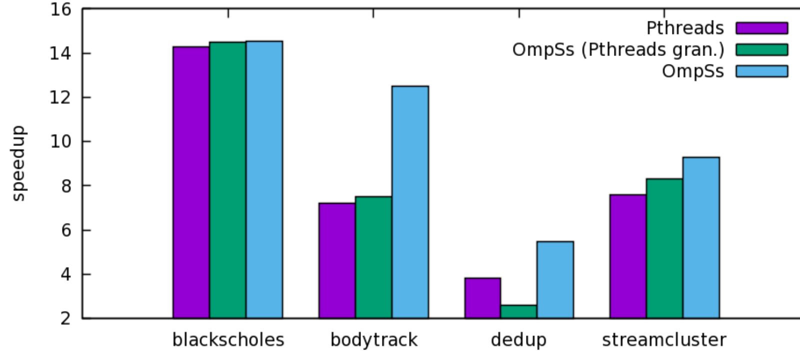

Fig. 7: Speedup comparison for Pthreads and OmpSs with the same granularity, as well as our optimized OmpSs version, when run on 16 cores. Results are normalized to the sequential version of the original code.

experiments on a two socket system. The Pthreads code partitions the working set by the number of available cores. We employ a different partition scheme to counter the NUMA effects. Through experimentation we observe that the best results can be obtained when using 80 blocks.

5.3. Task Granularity Impact

The granularity of individual tasks is an important factor that needs to be considered when parallelizing an application. Small task granularity can reduce load imbalance but such performance benefits can be neglected by the overhead of the runtime system, as it has to create and schedule more tasks. Results in Section 5.2 show how tuning the task granularity brings performance benefits in some cases (`blackscholes`, `bodytrack`, `dedup` and `streamcluster`) while in others it is better to keep the same parallel granularity as the PARSEC distribution codes (`canneal`, `facesim`, `ferret`, `freqmine` and `swaptions`).

In order to provide a more comprehensive comparison, this section examines the performance of `blackscholes`, `bodytrack`, `dedup` and `streamcluster` when using exactly the same granularity as in the PARSEC distribution code. Figure 7 shows the speedup of these benchmarks when run on 16 cores. The purple bar shows the speedup of the Pthreads version, the green one shows the speedup of a task-based implementation that has the same parallel granularity as its Pthreads counterpart. Finally, the light blue bar shows the speedup of the optimal task-based implementations discussed in Sections 3 and 5.2.

For the cases of `blackscholes` and `streamcluster` the parallelization scheme followed in the three codes (Pthreads and the two OmpSs versions) is the same. The difference between the two OmpSs versions is the granularity of the block sizes that are processed per task. In case of `blackscholes`, the OmpSs implementation with the same granularity as Pthreads does not perform better since the parallelism of this benchmark follows a fork-join model. In the case of `streamcluster`, the task-based implementations always improve the Pthreads performance, even if they operate following the same parallelization scheme and granularity as the Pthreads version. These improvements come from the NUMA effects correction that the OmpSs versions carry out.

In the case of `bodytrack` the optimal OmpSs implementation follows a quite different parallelization scheme than the original Pthreads code, as explained in Section 3. We

consider a trivial implementation in OmpSs where we follow the same parallelization strategy as in Pthreads. As shown in Figure 7, we do not observe any significant difference in performance among Pthreads and the equivalent OmpSs implementation. However, the new parallelization scheme is not applicable to Pthreads as it requires to synchronize the workload by explicit dependencies, which are not available in the Pthreads API.

In the case of `Dedup`, the trivial Pthreads-like implementation performs poorly, achieving a speedup of 2.6x. In this implementation, each pipeline stage is taskified following the Pthreads approach. Each large chunk is partitioned into smaller chunks, that will spawn three new tasks (`Compress`, `Deduplicate` and `WriteOutput`). This level of granularity creates hundreds of thousands of tasks, increasing the runtime's overhead significantly. In contrast, the optimized task-based version operates at the granularity of the large chunks, creating only a few hundreds of tasks, effectively reducing the runtime overhead.

In some cases, OmpSs can over-perform Pthreads even if the same parallelization scheme and granularity is followed, like the `streamcluster` results demonstrate. In some other cases (`dedup` and `bodytrack`), the performance improvements come from an optimized parallelization scheme. Such new schemes could be hardly implemented in Pthreads since they require a direct synchronization via explicit dependencies, which is not available in the Pthreads API. Finally, in case of simple fork-join applications (i. e. `blackscholes`) our performance benefits just come from further optimizing the parallel granularity.

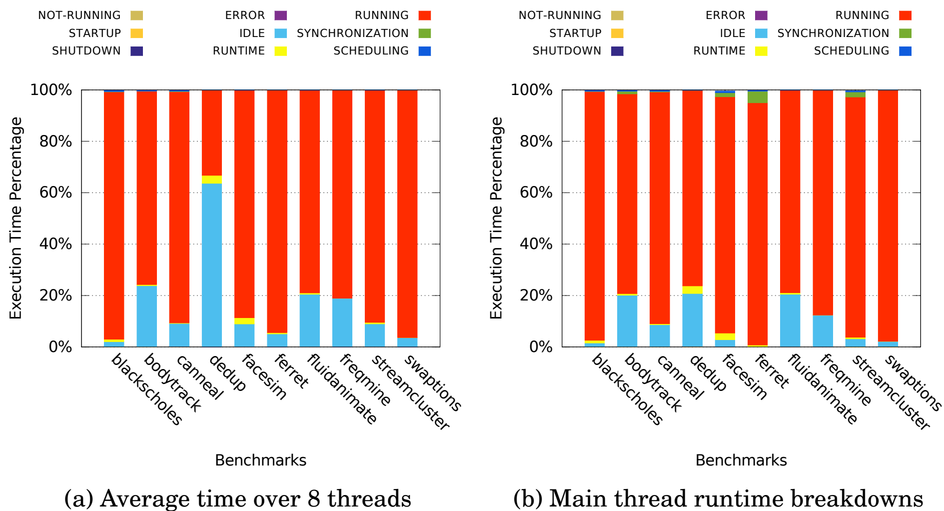

(a) Average time over 8 threads

(b) Main thread runtime breakdowns

Fig. 8: Runtime breakdowns when running on an 8-core configuration.

5.4. Runtime System Overhead

Task creation, scheduling and data dependencies tracking are all handled by the OmpSs runtime system. In this section, we evaluate the impact of these activities over the final parallel performance. Figure 8a shows a breakdown of the total execution time of each application. Each bar shows the breakdown of one application after averaging the values over eight concurrently executing threads. The red color represents the portion of time dedicated in running tasks, that is, in running user code. All

Table II: PARSEC parallelization model and properties characterization.

Benchmark	Parallel Model	I/O Heavy	Synchronization	LOC Reduction	Perf. Impr.
blackscholes	data-parallel	✗	dataflow	5.4%	0%
bodytrack	pipeline	✓	dataflow	81%	42%
canneal	unstructured	✗	locks/atomics	0%	-6.2%
dedup	pipeline	✓	dataflow/locks	38%	30%
facesim	pipeline	✗	dataflow/barrier	31%	34%
ferret	pipeline	✗	dataflow	46%	0%
fluidanimate	data-parallel	✗	dataflow/barrier	21%	5.7%
freqmine	data-parallel	✗	barrier/locks	0%	2.7%
streamcluster	data-parallel	✗	dataflow/barrier/atomics	33%	18%
swaptions	data-parallel	✗	dataflow	15%	6.6%

applications, excluding dedup, spend more than 75% of time doing useful work. The cyan bar represents idle time, which corresponds to the time a thread is waiting for some work to become available and is caused by load imbalance and sequential code phases. In most cases this time is low, except for `dedup`, where it reaches 60% of the total execution time. In Figure 8a we also represent the time spent in other activities like synchronization, scheduling, etc. None of these activities represent more than 5% of the total execution time.

Figure 8b shows the same breakdown of execution time but only for the main thread of execution, which is the one that runs serial parts of the code besides parallel tasks. `Dedup` has significantly lower idle time in the main thread, which indicates that there is not enough parallel work to keep all threads busy. This issue has been previously reported [Vandierendonck et al. 2013]. In general we see that the overhead of the runtime system is low, with only a few cases that show some time spent in synchronization, scheduling, and miscellaneous runtime overhead (in light yellow). Synchronization time can be time spent waiting a barrier or acquiring/releasing locks. Scheduling includes time needed to resolve dependencies and make scheduling decisions, while other runtime overheads are related to activity that cannot be associated with task scheduling and creation. Overall, we have seen that our implementations improve scalability considerably (by 13% on average), while runtime overhead remains low.

5.5. Characterization of the Applications

In Table II we characterize the considered applications in terms of parallelization model, I/O intensity and synchronization scheme. The table also shows code reductions and performance improvements achieved on a 16-core Sandy Bridge system. This table summarizes the properties of applications that make them good candidates for adopting a task-based model.

Applications characterized as data-parallel are limited to loop parallelism, where tasks are merely emulating an OpenMP loop construct. In these cases there is no performance gain, and the programming effort involved either with Pthreads, OpenMP or tasks, is similar. Pipeline applications are better candidates since they separate the application into discrete abstract stages. Implementing this paradigm with tasks implies taskifying the functionality of each stage and describing the data or control dependencies between them. In Pthreads, the programmer has to implement application specific thread pools and queuing systems to achieve the same performance. Also, task-based models offer in many cases an opportunity to easily expand the pipeline stages of the application with sequential and I/O intensive codes (e.g. `facesim` and `bodytrack` respectively). Indeed, by replacing locks and barriers, the runtime can discover additional dynamic parallelism and eliminate the cost of acquiring locks. Our task-based parallelization strategies successfully scale up the pipeline applications with a poorly scaling Pthread version (`bodytrack`, `dedup` and `facesim`) while reducing the code com-

plexity in all of them. In case of `ferret`, the task based version does not perform better than the Pthreads counterpart since its scalability is already very good (14x on a 16-core machine). The reduction in terms of lines of code is however dramatic: 46%. In the case of unstructured programs, e.g. `canneal`, task based programming does not offer any advantage over threading approaches.

Overall, we conclude that task-based parallelism can be effectively used to reduce the effort required to implement pipeline parallelism, while there are also important performance benefits to be gained if the application has no specific thread pooling mechanisms or I/0 intensive serial regions. In this scenario, the pipeline can be easily expanded to include the I/0 region and overlap it with a computation stage of the pipeline.

6. RELATED WORK

Few studies exist that examine the performance of task parallelism compared to other models. [Ayguadé et al. 2008] evaluate OpenMP tasks by implementing a few small kernel applications using the new OpenMP task construct. Their evaluation tests the model's expressiveness and flexibility as well as performance. [Podobas and Brorsson 2010] compare three models that implement task parallelism, Wool, Cilk++ and OpenMP. They compare their performance using small kernels, as well as some microbenchmarks aimed to measure task creation and synchronization costs. They show that Cilk++ and Wool have similar performance, while they outperform OpenMP tasks for fine grain workloads. On coarser grain loads, all models have matching performance with OpenMP gaining in one case, due to superior task scheduling.

BDDT [Tzenakis et al. 2012] is a task-based parallel model, very similar to OmpSs, that also uses a runtime to track data dependencies among tasks. BDDT uses block-level argument dependence tracking, where task arguments are processed into blocks of arbitrary size, which is defined by the user. BDDT is shown to outperform loop constructs implemented using OpenMP 2.0.

Other studies exist that compare parallel programming models in the literature. Although these studies do not focus on task parallelism, they employ benchmarks and similar methodology to evaluate their target models. [Coarfa et al. 2005] study and compare the performance of UPC and Co-array Fortran, two PGAS languages. They use select benchmarks from the NAS benchmark suite. [Appeltauer et al. 2009] use microbenchmarks to measure and compare the performance of 11 context-oriented languages. Their study shows that they all often manifest high overheads.

Although all the works we mention try to evaluate various programming models, in terms of performance, and some times on usability and versatility, they are all limited to small kernels or even just micro-benchmarks. We find that this approach is not sufficient to give us an insight on how a model will impact actual large-scale applications. [Karlin et al. 2013] use a proxy application in their work to evaluate a number of different programming models (OpenMP, MPI, MPI+OpenMP, CUDA, Chapel, Charm++, Liszt, Loci). Their approach however is limited to only one application. Different application domains can be very different, and may require different parallelization techniques to get good scalability and performance. A programming model could fail to even provide a way to express a parallelization scheme, let alone deliver performance. It is important to have an in depth understanding of a models behavior and limitation in order to make an educated decision whether research should direct its efforts to adopt and further expand it.

Pipeline parallelism has been the subject of study in some recent studies. This programming idiom is found often in streaming and server applications and goes far beyond the HPC domain. [Lee et al. 2013] propose an extension to the Cilk model, for expressing pipeline parallelism on-the-fly, without constructing the pipeline stages at

their dependencies a priori. It offers a performance comparison between the proposed model, Pthreads and Thread Building Blocks (TTB) for three PARSEC benchmarks, ferret, dedup and x264.

7. CONCLUSIONS

In this work we evaluate the benefits of task-based parallelisim by applying it to the PARSEC benchmark suite. We discuss and compare our implementations to their PARSEC Pthreads/OpenMP counterparts. We show how task parallelism can be applied on a wide range of applications from different domains. In fact, by comparing the lines of code between our implementations and the original versions, we make a strong case that task-based models are actually easier to use. The asynchronous nature of task-based parallelism, along with data dependence tracking through dataflow annotations, allows us to overlap computation with I/O phases. The underlying runtime system can take care of issues like scheduling and load balancing without significant overhead.

Our experimental results demonstrate that the task model can be easily applied on a wide range of applications beyond the HPC domain. Although, not all applications can benefit from a task-based approach, there are cases where it can greatly improve scalability. The programs that benefit most are those that present pipeline execution model, where different stages of the application can run concurrently. Finally, we plan to make a public release of our task-based implementations to stimulate research on these novel programming models.

8. ACKNOWLEDGEMENTS

This work has been partially supported by the European Research Council under the European Union 7th FP, ERC Grant Agreement number 321253, by the Spanish Ministry of Science and Innovation under grant TIN2012-34557, by the Severo Ochoa Program, awarded by the Spanish Government, under grant SEV-2011-00067 and by the HiPEAC Network of Excellence. M. Moreto has been partially supported by the Ministry of Economy and Competitiveness under Juan de la Cierva post-doctoral fellowship number JCI-2012-15047, and M. Casas is supported by the Secretary for Universities and Research of the Ministry of Economy and Knowledge of the Government of Catalonia and the Co-fund programme of the Marie Curie Actions of the 7th R&D Framework Programme of the European Union (Contract 2013 BP B 00243). Finally, the authors are grateful to the reviewers for their valuable comments, to the people from the Programming Models Group at BSC for their technical support, to the RoMoL team, and to Xavier Teruel, Roger Ferrer and Paul Caheny for their help in this work.

REFERENCES

Aaron B. Adcock, Blair D. Sullivan, Oscar R. Hernandez, and Michael W. Mahoney. 2013. Evaluating OpenMP Tasking at Scale for the Computation of Graph Hyperbolicity.. In *IWOMP (Lecture Notes in Computer Science)*, Vol. 8122. Springer, 71–83. DOI:http://dx.doi.org/10.1007/978-3-642-40698-0_6

Malte Appeltauer, Robert Hirschfeld, Michael Haupt, Jens Lincke, and Michael Perscheid. 2009. A Comparison of Context-oriented Programming Languages. In *International Workshop on Context-Oriented Programming (COP '09)*. ACM, Article 6, 6 pages. DOI:http://dx.doi.org/10.1145/1562112.1562118

Eduard Ayguadé, Nawal Copty, Alejandro Duran, Jay Hoeflinger, Yuan Lin, Federico Massaioli, Xavier Teruel, Priya Unnikrishnan, and Guansong Zhang. 2009. The Design of OpenMP Tasks. *IEEE Trans. Parallel Distrib. Syst.* 20, 3 (March 2009), 404–418. DOI:http://dx.doi.org/10.1109/TPDS.2008.105

Eduard Ayguadé, Alejandro Duran, Jay Hoeflinger, Federico Massaioli, and Xavier Teruel. 2008. Languages and Compilers for Parallel Computing. Springer-Verlag, Chapter An Experimental Evaluation of the New OpenMP Tasking Model, 63–77. DOI:http://dx.doi.org/10.1007/978-3-540-85261-2_5

Prithviraj Banerjee. 1994. *Parallel Algorithms for VLSI Computer-aided Design*. Prentice-Hall, Inc.

Pieter Bellens, Josep M. Perez, Rosa M. Badia, and Jesus Labarta. 2006. CellSs: a programming model for the cell BE architecture. In *SC*. Article 86. http://doi.acm.org/10.1145/1188455.1188546

Christian Bienia. 2011. *Benchmarking Modern Multiprocessors*. Ph.D. Dissertation. Princeton University.

Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques (PACT '08)*. ACM, 72–81. DOI:http://dx.doi.org/10.1145/1454115.1454128

Fischer Black and Myron S Scholes. 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81, 3 (May-June 1973), 637–54. http://ideas.repec.org/a/ucp/jpolec/v81y1973i3p637-54.html

Robert D. Blumofe, Christopher F. Joerg, Bradley C. Kuszmaul, Charles E. Leiserson, Keith H. Randall, and Yuli Zhou. 1995. Cilk: An Efficient Multithreaded Runtime System. *SIGPLAN Not.* 30, 8 (Aug. 1995), 207–216. DOI:http://dx.doi.org/10.1145/209937.209958

David R. Butenhof. 1997. *Programming with POSIX Threads*. Addison-Wesley Longman Publishing Co., Inc.

Barbara Chapman. 2007. The multicore programming challenge. In *Advanced Parallel Processing Technologies*. Springer. DOI:http://dx.doi.org/10.1007/978-3-540-76837-1_3

Barbara Chapman, Gabriele Jost, and Ruud van der Pas. 2007. *Using OpenMP: Portable Shared Memory Parallel Programming (Scientific and Engineering Computation)*. The MIT Press.

Cristian Coarfa, Yuri Dotsenko, John Mellor-Crummey, François Cantonnet, Tarek El-Ghazawi, Ashrujit Mohanti, Yiyi Yao, and Daniel Chavarría-Miranda. 2005. An Evaluation of Global Address Space Languages: Co-array Fortran and Unified Parallel C. In *Proceedings of the Tenth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '05)*. ACM, 36–47. DOI:http://dx.doi.org/10.1145/1065944.1065950

Henry Cook, Miquel Moreto, Sarah Bird, Khanh Dao, David A. Patterson, and Krste Asanovic. 2013. A Hardware Evaluation of Cache Partitioning to Improve Utilization and Energy-efficiency While Preserving Responsiveness. *SIGARCH Comput. Archit. News* 41, 3 (June 2013), 308–319. DOI:http://dx.doi.org/10.1145/2508148.2485949

Jack Dongarra, Robert Graybill, William Harrod, Robert F. Lucas, Ewing L. Lusk, Piotr Luszczek, Janice McMahon, Allan Snavely, Jeffrey S. Vetter, Katherine A. Yelick, Sadaf R. Alam, Roy L. Campbell, Laura Carrington, Tzu-Yi Chen, Omid Khalili, Jeremy S. Meredith, and Mustafa M. Tikir. 2008. DARPA's HPCS Program- History, Models, Tools, Languages. *Advances in Computers* 72 (2008), 1–100. DOI:http://dx.doi.org/10.1016/S0065-2458(08)00001-6

Alejandro Duran, Eduard Ayguadé, Rosa M. Badia, Jesús Labarta, Luis Martinell, Xavier Martorell, and Judit Planas. 2011. OmpSs: a Proposal for Programming Heterogeneous Multi-Core Architectures. *Parall. Proc. Lett.* 21, 2 (2011), 173–193. DOI:http://dx.doi.org/10.1007/978-3-642-37658-0_7

Alejandro Duran, Roger Ferrer, Eduard Ayguadé, RosaM. Badia, and Jesus Labarta. 2009. A Proposal to Extend the OpenMP Tasking Model with Dependent Tasks. *Int. J. Parallel Prog.* 37, 3 (2009), 292–305. DOI:http://dx.doi.org/10.1007/s10766-009-0101-1

Kayvon Fatahalian, Daniel Reiter Horn, Timothy J. Knight, Larkhoon Leem, Mike Houston, Ji Young Park, Mattan Erez, Manman Ren, Alex Aiken, William J. Dally, and Pat Hanrahan. 2006. Sequoia: Programming the Memory Hierarchy. In *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing (SC '06)*. ACM, Article 83. DOI:http://dx.doi.org/10.1145/1188455.1188543

Gsta Grahne and Jianfei Zhu. 2003. Efficiently Using Prefix-trees in Mining Frequent Itemsets.. In *FIMI* (2004-04-27) *(CEUR Workshop Proceedings)*, Bart Goethals and Mohammed Javeed Zaki (Eds.), Vol. 90. CEUR-WS.org. http://dblp.uni-trier.de/db/conf/fimi/fimi2003.html#GrahneZ03

Jiawei Han, Jian Pei, and Yiwen Yin. 2000. Mining Frequent Patterns Without Candidate Generation. *SIGMOD Rec.* 29, 2 (May 2000), 1–12. DOI:http://dx.doi.org/10.1145/335191.335372

David Heath, Robert Jarrow, and Andrew Morton. 1992. Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation. *Econometrica* 60, 1 (January 1992), 77–105. http://ideas.repec.org/a/ecm/emetrp/v60y1992i1p77-105.html

James Christopher Jenista, Yong hun Eom, and Brian Charles Demsky. 2011. OoOJava: Software Out-of-order Execution. *SIGPLAN Not.* 46, 8 (Feb. 2011), 57–68. DOI:http://dx.doi.org/10.1145/2038037.1941563

Ian Karlin, Abhinav Bhatele, Jeff Keasler, Bradford L. Chamberlain, Jonathan Cohen, Zachary Devito, Riyaz Haque, Dan Laney, Edward Luke, Felix Wang, David Richards, Martin Schulz, and Charles H. Still. 2013. Exploring Traditional and Emerging Parallel Programming Models Using a Proxy Application. In *Proceedings of the 2013 IEEE 27th International Symposium on Parallel and Distributed Processing (IPDPS '13)*. IEEE Computer Society, 919–932. DOI:http://dx.doi.org/10.1109/IPDPS.2013.115

I-Ting Angelina Lee, Charles E. Leiserson, Tao B. Schardl, Jim Sukha, and Zhunping Zhang. 2013. On-the-fly Pipeline Parallelism. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '13)*. ACM, New York, NY, USA, 140–151. http://doi.acm.org/10.1145/2486159.2486174

Qin Lv, William Josephson, Zhe Wang, Moses Charikar, and Kai Li. 2006. Ferret: A Toolkit for Content-based Similarity Search of Feature-rich Data. *SIGOPS Oper. Syst. Rev.* 40, 4 (April 2006), 317–330. DOI:http://dx.doi.org/10.1145/1218063.1217966

Matthias Müller, David Charypar, and Markus Gross. 2003. Particle-based Fluid Simulation for Interactive Applications. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '03)*. Eurographics Association, 154–159. http://dl.acm.org/citation.cfm?id=846276.846298

Dan Nagle. 2005. MPI – The Complete Reference, Vol. 1, The MPI Core, 2Nd Ed., Scientific and Engineering Computation Series, by Marc Snir, Steve Otto, Steven Huss-Lederman, David Walker and Jack Dongarra. *Sci. Program.* 13, 1 (Jan. 2005), 57–63. DOI:http://dx.doi.org/10.1155/2005/653765

Judit Planas, Rosa M. Badia, Eduard Ayguadé, and Jesus Labarta. 2009. Hierarchical Task-Based Programming With StarSs. *Int. J. High Perform. Comput. Appl.* 23, 3 (Aug. 2009), 284–299. DOI:http://dx.doi.org/10.1177/1094342009106195

Artur Podobas and Mats Brorsson. 2010. A Comparison of some recent Task-based Parallel Programming Models. In *Proceedings of the 3rd Workshop on Programmability Issues for Multi-Core Computers, (MULTIPROG'2010), Jan 2010, Pisa*. Qc 20120214.

Sean Quinlan and Sean Dorward. 2002. Awarded Best Paper! - Venti: A New Approach to Archival Data Storage. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies (FAST '02)*. USENIX Association, Article 7. http://dl.acm.org/citation.cfm?id=1083323.1083333

Eftychios Sifakis, Igor Neverov, and Ronald Fedkiw. 2005. Automatic Determination of Facial Muscle Activations from Sparse Motion Capture Marker Data. *ACM Trans. Graph.* 24, 3 (July 2005), 417–425. DOI:http://dx.doi.org/10.1145/1073204.1073208

Christi Symeonidou, Polyvios Pratikakis, Angelos Bilas, and Dimitrios S. Nikolopoulos. 2013. DRASync: Distributed Region-based Memory Allocation and Synchronization. In *Proceedings of the 20th European MPI Users' Group Meeting (EuroMPI '13)*. ACM, 49–54. DOI:http://dx.doi.org/10.1145/2488551.2488558

George Tzenakis, Angelos Papatriantafyllou, John Kesapides, Polyvios Pratikakis, Hans Vandierendonck, and Dimitrios S. Nikolopoulos. 2012. BDDT:: Block-level Dynamic Dependence Analysisfor Deterministic Task-based Parallelism. *SIGPLAN Not.* 47, 8 (Feb. 2012), 301–302. DOI:http://dx.doi.org/10.1145/2370036.2145864

Hans Vandierendonck, Kallia Chronaki, and Dimitrios S. Nikolopoulos. 2013. Deterministic Scale-free Pipeline Parallelism with Hyperqueues. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC '13)*. ACM, Article 32, 12 pages. http://doi.acm.org/10.1145/2503210.2503233

Hans Vandierendonck, Polyvios Pratikakis, and Dimitrios S. Nikolopoulos. 2011. Parallel Programming of General-purpose Programs Using Task-based Programming Models. In *Proceedings of the 3rd USENIX Conference on Hot Topic in Parallelism (HotPar'11)*. USENIX Association, 13–13. http://dl.acm.org/citation.cfm?id=2001252.2001265

Stéphane Zuckerman, Joshua Suetterlein, Rob Knauerhase, and Guang R. Gao. 2011. Using a "Codelet" Program Execution Model for Exascale Machines: Position Paper. In *EXADAPT '11*. ACM, 64–69. http://doi.acm.org/10.1145/2000417.2000424