Provably and Practically Efficient Granularity Control

Umut A. Acar
Carnegie Mellon University and Inria
USA and France
umut@cs.cmu.edu

Arthur Charguéraud
Inria & Université de Strasbourg, CNRS, ICube
France
arthur.chargueraud@inria.fr

Vitaly Aksenov
Inria and ITMO University
France and Russia
vitalii.aksenov@inria.fr

Mike Rainey
Indiana University and Inria
USA and France
me@mike-rainey.site

Abstract

Over the past decade, many programming languages and systems for parallel-computing have been developed, e.g., Fork/Join and Habanero Java, Parallel Haskell, Parallel ML, and X10. Although these systems raise the level of abstraction for writing parallel codes, performance continues to require labor-intensive optimizations for coarsening the granularity of parallel executions. In this paper, we present provably and practically efficient techniques for controlling granularity within the run-time system of the language. Our starting point is “oracle-guided scheduling”, a result from the functional-programming community that shows that granularity can be controlled by an “oracle” that can predict the execution time of parallel codes. We give an algorithm for implementing such an oracle and prove that it has the desired theoretical properties under the nested-parallel programming model. We implement the oracle in C++ by extending Cilk and evaluate its practical performance. The results show that our techniques can essentially eliminate hand tuning while closely matching the performance of hand tuned codes.

CCS Concepts • Software and its engineering → Parallel programming languages;

Keywords parallel programming languages, granularity control

1 Introduction

The proliferation of multicore hardware in the past decade has brought shared-memory parallelism into the mainstream.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PPoPP '19, February 16–20, 2019, Washington, DC, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6225-2/19/02...\$15.00

<https://doi.org/10.1145/3293883.3295725>

This change has led to much research on *implicit threading*, or, alternatively, implicit parallelism, which goes back to early parallel programming languages such as Id [Arvind and Gostelow 1978] and Multilisp [Halstead 1985]. Implicit threading seeks to make parallel programming easier by delegating certain tedious but important details, such as the scheduling of parallel tasks, to the compiler and the run-time system. There has been much work on specialized programming languages and language extensions for parallel systems, for various forms of implicit threading, including OpenMP [OpenMP Architecture Review Board 2008], Cilk [Frigo et al. 1998], Fork/Join Java [Lea 2000], Habanero Java [Imam and Sarkar 2014], NESL [Blelloch et al. 1994], TPL [Leijen et al. 2009], TBB [Intel 2011], X10 [Charles et al. 2005], parallel ML [Fluet et al. 2011; Guatto et al. 2018; Jagannathan et al. 2010; Raghunathan et al. 2016], and parallel Haskell [Chakravarty et al. 2007; Keller et al. 2010].

Nearly all of these languages support a lightweight syntax for expressing parallelism. For example, in Cilk Plus, the programmer needs two keywords to express parallelism, **spawn** and **sync**, which respectively indicate a computation that can be executed in parallel and a computation that must be synchronized with. These two keywords suffice to express parallel loops and *nested parallelism*, wherein parallel computations may themselves start and synchronize with other parallel computations. When combined with the powerful features of modern programming languages, these simple constructs enable elegant expression of parallelism. For example, we can implement a “parallel map” that applies some function f to each element of an array in a C++-like parallel language by using a parallel-for loop (`pfor`), as shown in the `map` function in Figure 1. By using templates, this implementation enables many different forms of mapping, e.g., `map` can be used to increment the elements in an input array of integers or it can be used to implement a matrix-vector multiplication by passing as an argument an array of arrays (vectors) representing the matrix.

Even though implicit parallelism can be elegant, and offer important asymptotic efficiency guarantees [Acar et al. 2002, 2013; Blelloch and Greiner 1996; Blumofe and Leiserson 1999; Frigo et al. 1998], there remains a key concern: the constants

```

template <F,T,S>
void map(F f, T* a, S* b, int n)
    pfor (i = 0; i < n; i++)
        b[i] = f(a[i])

int grain = ... // to be determined
template <F,T,S>
void map_coarsened(F f, T* a, S* b, int n)
    pfor (i = 0; i < (n + grain - 1) / grain; i++)
        for (j = i * grain; j < min(n, (i + 1) * grain); j++)
            b[j] = f(a[j]);

```

Figure 1. Parallel map, uncoarsened and coarsened implementations.

hidden in these asymptotic analyses, a.k.a., the *overheads of parallelism* can easily overwhelm the benefits of parallelism in practice. This challenge is especially difficult in high-level, declarative codes, such as map, and is sometimes considered a “holy grail” in parallelism research (e.g., [Tzannes et al. 2014]). For example, the simple map example above could run as much as 10-100x slower than an optimized implementation due to these overheads.

To control the overheads of parallelism, the current state of the art requires the programmer to *tune* the code to perform *granularity control* or *coarsening* so as to amortize the overheads [Intel 2011]. To do so, the programmer identifies for each potential parallel computation a *sequential alternative*, a semantically equivalent sequential piece of code, and makes sure that this sequential alternative is executed for small (and only for small) computations. As an example, imagine an instance of map where the programmer increments each element of an integer array by one. To perform this operation efficiently, the programmer should “bunch” iterations into blocks that are executed sequentially and select the bunch size to be large enough to amortize the cost of parallelism. The function map_coarsened in Figure 1 shows such a tuned implementation of map, where the value grain determines the bunch size.

The optimized map implementation is far from its elegant predecessor, but elegance aside, it is not easy to make map perform well, because the programmer still has to choose the setting for grain. There are at least two problems with this approach.

- The optimal setting for the grain depends on the architecture as well as the specific inputs and requires a manual, labor-intensive search [Bergstrom et al. 2012; Feeley 1992, 1993a; Hiraishi et al. 2009; Intel 2011; Lopez et al. 1996; Rainey 2010; Sanchez et al. 2010a; Tzannes et al. 2014].
- As noted by many researchers, such optimizations could also harm performance portability by overfitting the code to a particular machine [Acar et al. 2016;

Huelsbergen et al. 1994; Lopez et al. 1996; Pehoushek and Weening 1990; Tzannes et al. 2014; Weening 1989].

- It can be impractical to tune generic codes such as map, because the optimal grain depends on the template parameters as well as the actual arguments to map, some of which, such as the function argument, might not be known a priori.

Prior work therefore proposed techniques for (semi-) automating this process [Huelsbergen et al. 1994; Lopez et al. 1996; Pehoushek and Weening 1990; Weening 1989]. Yet, these techniques remain heuristics that fall short of delivering strong theoretical and practical efficiency guarantees. A more recent paper on oracle-guided scheduling [Acar et al. 2016] shows that strong bounds can be achieved, but assumes an oracle that can predict the future. The authors do present an implementation for an oracle, but their implementation technique can support only certain flat-parallel programs, where parallel computations cannot be nested in other parallel computations. As such, this limitation excludes nearly all interesting parallel programs expressible in high-level parallel programs today.

In this paper, we present provably efficient and practical techniques for controlling granularity for nested-parallel programs. As in oracle-guided scheduling [Acar et al. 2016], we ask the programmer to provide an *abstract cost function*, e.g., asymptotic cost, for each piece of parallel code. We then provide an online prediction algorithm that uses such information to implement the oracle by predicting the actual work (1-processor run-time) of parallel codes. Our major contribution is this prediction algorithm and the result that the algorithm delivers efficient and practical granularity control, for any nested-parallel program, under a certain number of assumptions that we present in detail. The key insight behind our algorithm is to break a circular dependency encountered in implementing the oracle in prior work [Acar et al. 2016] by using an iterative refinement algorithm that, after a small number of iterations, provably converges to the desired values. We present a theoretical analysis of the key efficiency properties of the algorithm and evaluate the approach by comparison to a highly (hand) optimized suite of benchmarks from the Problem Based Benchmark Suite [Blelloch et al. 2012]. Our results show that our techniques can eliminate the need for hand tuning in many cases and closely match the performance of the hand optimized code.

The specific contributions of this paper include:

- an online algorithm for automatic granularity control in lightly annotated, nested-parallel computations,
- end-to-end tight bounds on run time or parallel codes,
- an implementation for Cilk Plus, which implements work stealing for C++, and
- an evaluation on a broad collection of hand-tuned benchmarks in Cilk Plus.

A long version of the paper is also available [Acar et al. 2019].

```

T f(x) {
  T r // local variable to store the result
  spguard([&] { return c(x) },           // cost function
    [&] { if |x| == 1 ...                // parallel body
      else (x1, x2) = divide(x)
        r1 = spawn f(x1)
        r2 =      f(x2)
        sync
        r = conquer(r1, r2) },
    [&] { r = g(x) })                  // sequential body
  return r }

```

Figure 2. A sample use of `spguard`. Higher-order functions are expressed using the C++ lambda-expression syntax.

2 Algorithmic granularity control

Our goal is to transfer the burden of granularity control to a capable, library implementation. To this end, we ask the programmer to provide for each parallel function a *series-parallel guard*, by using the keyword `spguard`. A `spguard` consists of: a *parallel body*, which is a lambda function that performs a programmer-specified parallel computation; a *sequential body*, which is a lambda function that performs a purely sequential computation equivalent to the parallel body, i.e., performing the same side-effects and delivering the same result, a *cost function*, which gives an abstract measure, as a positive number, of the work (run-time cost) that would be performed by the sequential body.

At a high level, a `spguard` exploits the result of the cost function to determine whether the computation involved is small enough to be executed sequentially, i.e., without attempting to spawn any subcomputation. If so, the `spguard` executes the sequential body. Otherwise, it executes the parallel body, which would typically spawn smaller subcomputations, each of them being similarly guarded by a `spguard`.

To see the basic idea behind our algorithm, consider the example shown in Figure 2, involving a single `spguard`. Suppose that we have a parallel function $f(x)$ which, given an argument x , performs some computation in divide-and-conquer fashion, by recursively calling itself in parallel, as shown above. Assume the body of this function to be controlled by a `spguard`, with $c(x)$ denoting the cost function, and $g(x)$ denoting the sequential body, that is, a purely sequential function that computes the same result as $f(x)$.

The cost function may be any programmer-specified piece of code that, given the context, computes a value in proportion to the one-processor execution time of the sequential body. Typically, the cost function depends on the arguments provided to the current function call. A good choice for the cost function is the average asymptotic complexity of the sequential body, e.g., $n \log n$, or n , or $\sqrt{n}$, where n denotes the size of the input. The programmer need not worry about constant factors because `spguards` are able to infer them on-line, with sufficient accuracy.

In a real implementation, the sequential body can be left implicit in many cases, because it can be inferred automatically. For example, the sequential body for a parallel-for loop can be obtained by replacing the parallel-for primitive with a sequential-for. Likewise, in many instances, the complexity function is linear, allowing us to set it to the default when not specified. In our library and experiments, we use this approach to dramatically reduce the annotations needed.

2.1 Sequentialization decisions

Our algorithm aims at sequentializing computations that involve no more than a small amount of work. To quantify this amount, let us introduce the *parallelism unit*, written κ , to denote the smallest amount of work (in units of time) that would be profitable to parallelize on the host architecture. The value of κ should be just large enough to amortize the cost of creating and managing a single parallel task. On modern computers, this cost can be from hundreds to thousands of cycles. Practical values for κ therefore range between 25 and 500 microseconds.

Intuitively, we aim at enforcing the following policy: if the result of $f(x)$ can be obtained by evaluating the sequential body $g(x)$ in time less than κ , then $g(x)$ should be used. Under a small number of assumptions detailed in further, this policy leads to provably efficient granularity control.

One central question is how to predict whether a call to $g(x)$ would take less than κ units of time. Assume, to begin with, a favorable environment where (1) the hardware is predictable in the sense that the execution time is in proportion to the number of instructions, and (2) the cost function $c(x)$ gives an estimate of the asymptotic number of instructions involved in the execution of $g(x)$. For a given input x , let N denote the value of $c(x)$, and let T denote the execution time of $g(x)$. By definition of “asymptotic”, there exists a constant C such that: $T \approx C \cdot N$. Our algorithm aims at computing C by sampling executions, and then it exploits this constant to predict whether particular calls to $g(\cdot)$ take fewer than κ units of time.

More precisely, for an input x_i being evaluated sequentially, that is, through a call to $g(x_i)$, we may measure the execution time of $g(x_i)$, written T_i , and we may compute the value of $c(x_i)$, written N_i . From a collection of samples of the form (T_i, N_i) , we may evaluate the value of the constant C by computing the ratios T_i/N_i . In reality, the actual value of the constant can vary dramatically depending on the size of the computation, in particular due to cache effects—we will return to that point. There is a much bigger catch to be addressed first.

In order to decide which computations are safe to execute using the sequential body $g(\cdot)$, our algorithm needs to first know the constant C . Indeed, without a sufficiently accurate estimate of the constant, the algorithm might end up invoking $g(\cdot)$ on a large input, thereby potentially destroying all available parallelism. Yet, at the same time, in order

to estimate the constant C , the algorithm needs to measure the execution time of invocations of $g(\cdot)$. Thus, determining the value of C and executing the algorithm $g(\cdot)$ are interdependent. Resolving this critical circular dependency is a key technical challenge.

At a high level, our algorithm progressively sequentializes larger and larger computations. It begins by sequentializing only the base case, and ultimately converges to computations of duration κ . Each time that we sequentialize a computation by calling $g(\cdot)$ instead of $f(\cdot)$, we obtain a measure. This measure may be subsequently used to predict that another, slightly larger input may also be processed sequentially. We are careful to increase the input size progressively, in order to always remain on the safe side, making sure that our algorithm never executes sequentially a computation significantly longer than κ .

Because our algorithm increases the cost (as measured by the cost function) of sequentialized computations each time at least by a multiplicative factor, called α , the estimation converges after just a logarithmic number of steps. The *growth rate*, α controls how fast sequentialized computations are allowed to grow. Any $\alpha > 1$ could be used; values between 1.2 and 5 work well in practice.

2.2 Nested parallelism

When dealing with a single spguard, the process described above generally suffices to infer the constant associated with that spguard. However, the process falls short for programs involving nested parallelism, e.g., nested loops or mutually-recursive functions. To see why, consider a function $h(\cdot)$ that consists of a spguard whose parallel body performs some local processing and then spawns a number of calls to completely independent functions. Because $h(\cdot)$ is not a recursive function with a base case, the algorithm described so far does not have a chance to follow the convergence process; the spguard of $h(\cdot)$ would have no information whatsoever about its constant, and it would always invoke the parallel body, consequently failing to control granularity.

To address this issue and support the general case of nested parallelism, we introduce an additional mechanism. When executing the parallel body of a spguard, our algorithm computes the sum of the durations of all the pieces of sequential computation involved in the execution of that parallel body. This value gives an upper bound on the time that the sequential body would have taken to execute. This upper bound enables deriving an over-approximation of the constant. Our algorithm uses this mechanism to make *safe* sequentialization decisions, i.e., to sequentialize computations that certainly require less than κ time. By measuring the duration of such sequential runs, our algorithm is then able to refine its estimate of the constant. It may subsequently sequentialize computations of size closer to κ .

Overall, our algorithm still progressively sequentializes larger and larger subcomputations, only it is able to do so by traversing distinct spguards with different constant factors.

2.3 Dealing with real hardware

Our discussion so far assumes that execution times may be predicted by the relationship $T \approx C \cdot N$. But in reality, this assumption is not the case. The reason is that the ratios T/N may significantly depend on the input size. For example, one can observe on a microbenchmark that processing an input that does not fit into the L3 cache may take up to 10 times longer to execute than a just slightly smaller input that fits in the cache. In fact, even two calls to the same function on the same input may have measured time several folds apart, for example, if the first call needs to load the data into the cache but not the second one.

We design our algorithm to be robust in the face of large variations of the execution times typical of modern hardware. In addition to validating empirically that our algorithm behaves well on hundreds of runs of our full benchmark suite, we formally prove its robustness property. To that end, we consider a relatively realistic model that takes into account the variability of execution times typical of current hardware. More precisely, we develop our theory with respect to an abstract notion of “work”, which we connect both to runtime measures and to results of costs functions, as described next.

First, we assume that runtime measures may vary by no more than a multiplicative factor E from the work, in either direction. Second, we require that, for the program and the growth rate α considered, there exists a value β such that, for any cost function involved in the program, the following property holds: if the cost function applied to input J returns a value no more than α bigger than for input I , then the work associated with input J is at most β times bigger than the work for input I . Intuitively, this property ensures that the work increases with the cost, but not exceeding a maximal rate of increase.

One important feature of our algorithm is that its execution does not require knowledge of E or β . These parameters are only involved in the analysis. It may nevertheless be useful to estimate the values of E and β that apply on the target hardware. In the long version of the paper, we show the results of estimating these parameters on our test harness.

2.4 Analysis

Our algorithm relies on several important assumptions. These assumptions, whose formal statements may be found in the long version of the paper, are matched by a large class of realistic parallel algorithms.

First, our algorithm assumes that, for each spguard, the sequential body evaluates no slower than the corresponding parallel body on a single processor. (The sequential body might always be obtained by replacing spawns with sequences in the parallel body.) Furthermore, the sequential

body should not run arbitrarily faster than the parallel body when executed with all of its inner spguards sequentialized. This assumption is required to know that it is safe to exploit the execution time of the parallel body to over-approximate that of the sequential body.

To better understand this latter part of the assumption, consider the following example. Assume a purely sequential algorithm, whose cost is $a \cdot N$ for an input of size N , for some constant a . Assume also a divide-and-conquer parallel algorithm that recursively divides its input in halves, such that for an input of size N , the division process (and, symmetrically, the conquer process) takes time $b \cdot N$, for some constant b ; the cost of this parallel algorithm is thus $b \cdot N \log N$. On the one hand, the sequential algorithm may be assumed to be faster than the parallel algorithm, because it is always possible to instantiate the sequential algorithm as the sequential execution of the parallel algorithm. On the other hand, let us check the requirement that “the sequential body does not run arbitrarily faster than the parallel body executed with all its inner spguards being sequentialized.” Executing one step of the divide-and-conquer process (as described by the parallel body) and then immediately calling the sequential body in both branches has cost: $b \cdot N + 2 \cdot a \cdot \frac{N}{2}$. Executing the sequential body directly has cost: $a \cdot N$. The ratio between these two costs is equal to: $1 + \frac{b}{a}$, which is bounded by a constant, as required. Many practical algorithms exhibit a similar pattern.

Second, our algorithm assumes that spguards evaluate in constant time and are called regularly enough through the call tree. Without this latter assumption, the program could have a spguard called on an input that takes much longer than κ to process, with the immediately nested spguard called on an input that takes much less than κ , leaving no opportunity for our algorithm to sequentialize computations of duration close to κ .

Third and last, our algorithm assumes some balance between the branches of a fork-join. Without this assumption, the program could be a right-leaning tree, with all left branches containing tiny computations. Such an ill-balanced program is a poorly-designed parallel program that would be slow in any case.

As we prove through a careful analysis detailed in the long version of the paper, under the aforementioned assumptions, our algorithm is efficient. Our bound generalizes Brent’s bound ($T_P \leq \frac{w}{P} + s$) [Zadeh 2017] by taking into account the overheads of thread creation (cost of spawning and managing threads) and granularity control (including the cost of evaluating cost functions, which are assumed to execute in constant time). A simplified statement of the bound, using big-O notation, appears next.

Theorem 2.1. *Under the above assumptions, with parallelism unit κ , the runtime on P processors using any greedy scheduler*

of a program involving work w and span s is bounded by:

$$T_P \leq \left(1 + \frac{O(1)}{\kappa}\right) \cdot \frac{w}{P} + O(\kappa) \cdot s + O(\log^2 \kappa).$$

The most important term in this bound is the first term: the overheads impact the work term w by only a small factor $O(1)/\kappa$, which can be reduced by controlling κ . The second term states that doing so increases the span by a small factor $O(\kappa)$, and that the granularity control comes at a small logarithmic overhead $O(\log^2 \kappa)$, which is due to our granularity control algorithm. Although our theorem is stated, for simplicity, in terms of greedy schedulers, the bound could be easily adapted, e.g., to work-stealing schedulers.

2.5 Pseudo-code for the estimator and spguard

In what follows, we present pseudo-code for the core of our algorithm. For each syntactic instance of a spguard, we require a unique *estimator* data structure to store the information required to estimate the corresponding constant factor. (In the case of templated code, we instrument the template system to ensure that each template instantiation allocates one independent estimator data structure.) The top of Figure 3 shows the code of the estimator data structure. The estimator data structure maintains only two variables: the variable C , a running estimate of the constant, and $N_{\max}$, which tracks the maximum (abstract) work sampled so far.

The function `report` provides the estimator with samples. It takes as argument T , the execution time, and N , the abstract work. If T is less than the parallelism threshold κ and N is greater than $N_{\max}$, then C is atomically updated as T/N and $N_{\max}$ is updated to N . The function `report` is protected against data races by using an atomic block. Our library implements this atomic block by using a single compare-and-swap (CAS) operation and packing the variables C (represented in single-precision) and $N_{\max}$ (represented on 32 bits) on a single machine word. In the case of a race, compare-and-swap would fail and our code would try again until it succeeds, or $N_{\max}$ becomes greater than N .

The function `is_small`, takes as argument a cost N and returns a boolean indicating whether the time to execute this much work is predicted to be less than $\alpha \cdot \kappa$. (Allowing sequentialization of computations of size slightly larger than κ is needed to obtain proper estimates near κ .) If N is smaller than $N_{\max}$, then the function returns true. Intuitively, this computation is smaller than a previously-sampled computation that took less than κ time. If N exceeds $N_{\max}$, but no more than by a factor α , the function extrapolates using its most current information: if the product $C \cdot N$ does not exceed $\alpha \cdot \kappa$, then it returns true.

The bottom of Figure 3 shows the functions `fork2join` and `sguard`. The `fork2join` function executes its two branches in parallel and resumes after completion of both branches. The function `sguard` takes an estimator, a parallel body,

```

1  const double κ    // parallelism unit
2  const double α    // growth factor
3
4  class estimator
5      double C      // constant for estimations
6      int Nmax = 0  // max complexity measure
7      void report(int N, time T)
8          atomic { if T ≤ κ and N > Nmax
9                  C = T / N
10                 Nmax = N }
11
12     bool is_small(int N)
13         return (N ≤ Nmax) or
14                (N ≤ α · Nmax and N · C ≤ α · κ)
15
16     template <Body_left, Body_right>
17     void fork2join(Body_left bl, Body_right br)
18         spawn bl()
19         br()
20     sync
21
22     template <Complexity, Par_body, Seq_body>
23     void spguard(estimator* es, Complexity c,
24                 Par_body pb, Seq_body sb)
25         int N = c()
26         time work = if es.is_small(N)
27                     then measured_run(sb)
28                     else measured_run(pb)
29     es.report(N, work)

```

Figure 3. Pseudocode for the core algorithm.

a sequential body and a cost function that return the abstract work of the sequential body. It begins by computing the abstract cost N for the sequential body and consults the estimator. If the work is predicted to be small, it runs the sequential body, else the parallel body. It relies on a function called `measured_run` to measure the sum of the duration of the pieces of sequential code executed. This function is not entirely trivial to implement, but can be implemented efficiently with little code (see the long version of the paper).

2.6 Robustness with respect to outliers

One crucial aspect of the design of our algorithm is its robustness. As explained in Section 2.3, a given run might be much slower than others, e.g., due to cache misses or OS interference. If such slow runs were taken into account for making future predictions, it could result in the sequentialization of tiny computations running in much less than κ time (under favorable conditions). Our algorithm deals with outlier measurements by keeping track of $N_{\max}$, which denotes the maximal N (i.e., result of the complexity function) for which a sequential run took less than κ time. Upon encountering a run that was much slower than previous runs,

the corresponding time measure exceeds κ time and thus gets discarded by our algorithm (Figure 3, Line 8).

Observe that our algorithm only ever increases $N_{\max}$. As we observed through experiments, there is an inherent asymmetry in the distribution of execution times, featuring a long tail. On the one hand, an execution time can be abnormally slow by a very large factor, due to, e.g., repeated OS interference or saturated memory bus. Our algorithm is carefully designed to be robust to such outliers, which we do observe in practice. On the other hand, an execution time can be abnormally fast only by some constant factor. It is possible that a warmed-up cache, or a data set unusually laid out in memory, could speed up the execution, but not arbitrarily. Our theoretical analysis accounts for that possibility through the error factor E (Section 2.3). As we prove, in the worst case, a run that would be abnormally fast by a factor E could only increase the total amount of overheads by a factor E .

We highlight two interesting features of our C++ implementation. First, to avoid always having to pass a sequential body explicitly, the `spguard` provides a default setting. If the `spguard` is called with such a default setting, then the parallel body is called whenever the `spguard` switches to sequential execution. However, in this scenario, when it sequentializes, the `spguard` executes the parallel body in an environment in which all `fork2join` calls are performed sequentially, using dynamic checks to selectively disable parallelism. Second, our implementation supports defining higher-level abstractions on top of `fork2join`, such as `parallel-for`, `map`, `reduce`, `map_reduce`, `scan`, `filter`, etc. For these abstractions, the programmer may either indicate a custom cost function or simply rely on the default one, which assumes a constant-time processing for each item or iteration. Because it applies often, the default cost function greatly reduces the number of hand-written cost functions.

2.7 Intuition for the proof

Without loss of expressiveness, we assume for the proof that the parallel body of a `spguard` consists of a sequence of (1) a sequential preprocessing, (2) a fork-join operation with left and right branches, and (3) a sequential postprocessing. (In practice, this amounts to surrounding every block of source code featuring a fork-join with a `spguard`.)

In that setting, we assume that the execution of a `spguard` (including the evaluation of its cost function) is assumed to incur a cost ϕ , and that each execution of a `spguard` that evaluates in parallel (i.e., that triggers a fork-join) is assumed to incur an extra cost τ . In other words, ϕ quantifies the overheads of granularity control, and τ quantifies the overheads of parallelism.

We define the *total work*, written $\mathbb{W}$, and the *total span*, written $\mathbb{S}$, to measure the actual work and span of an execution, i.e., accounting for the costs τ and ϕ . To derive our final theorem, we bound the total work and span ($\mathbb{W}$ and $\mathbb{S}$) in terms of the *algorithmic* work and span (w and s), which

excludes overheads. We then invoke Brent’s theorem [Zadeh 2017]: $T_p \leq \frac{\mathbb{W}}{p} + \mathbb{S}$, to bound the parallel execution time.

For the span, we establish that the total span can only be a multiplicative factor larger than the algorithmic span, in the sense that $\mathbb{S}$ is $O(\kappa) \cdot s$. More precisely, we establish the following bound by induction on the program execution:

$$\mathbb{S} \leq (1 + \phi + \max(\tau, E\beta\kappa)) \cdot s.$$

Recall that κ denotes the parallelism unit (Section 2.1), and that E and β are constant parameters introduced for the analysis (Section 2.3). The proof relies on a key lemma, which asserts that if a spguard executes its sequential body, then this body involves at most $E\beta\kappa$ sequential work. Intuitively, this lemma captures the fact that work predictions performed by our algorithm cannot be under-estimated by more than a factor $E\beta$. (We cannot expect tighter bounds in our model, for the reasons discussed in Section 2.3.)

The intuition for the bound $(1 + \phi + \max(\tau, E\beta\kappa)) \cdot s$ is as follows. At a spguard, the algorithmic span increases by one unit. If the corresponding spguard executes its parallel body, then its total span increases by $1 + \phi + \tau$, to account for overheads. If, however, the spguard executes its sequential body, then, because the work involved is at most $E\beta\kappa$, the total span increases at most by $1 + \phi + E\beta\kappa$. The bound is obtained by considering the maximum of these two cases.

For the work, we prove that the total work is only a fraction larger than the algorithmic work, plus a small constant:

$$\mathbb{W} \leq \left(1 + \frac{O(1)}{\kappa}\right) \cdot w + O(P \log^2 \kappa).$$

To establish this bound, we distinguish the spguard calls that involve $\Theta(\kappa)$ sequential work from those that involve less.

After the algorithm has converged, spguards are able to make predictions sufficiently accurate to ensure that the sequentialized spguards involve at least $\Theta(\kappa)$ sequential work.¹ When all sequentialized spguards involve $\Theta(\kappa)$ work, the constant time overheads associated with each spguard only amount for a fraction $\frac{O(1)}{\kappa}$ of the algorithmic work involved. For example, consider a simple parallel program performing binary division of its input. The execution of this program can be viewed as a binary tree. Let n denote the number of leaves in the tree. Every non-leaf node induces an overhead $\phi + \tau$, to pay for the evaluation of the spguard and the fork. Every node (including leaves) induces an overhead ϕ , to pay for the spguard. Because there are $n - 1$ nodes in a binary tree with n leaves, the overheads amount to $n \cdot \phi + (n - 1) \cdot (\phi + \tau)$, which is $O(n \cdot (\tau + 2\phi))$. These overheads can be amortized over the sequential work performed at the n leaves, each of

¹Technically, any spguard that involves less than $\kappa/(\gamma DE)$ work will necessarily execute its sequential body, where E denotes the error factor, where D bounds the ratio the work of the sequential body and the sequential work of the parallel body of a spguard (first assumption from Section 2.4), and where γ bounds the ratio between the sequential work of a spguard and that of the immediately nested spguard (second assumption from Section 2.4). Formal definitions for D and γ are given in the long version of the paper.

which involves $\Theta(\kappa)$ work. More precisely, the ratio between the overheads and the sequential work is $O\left(\frac{n \cdot (\tau + 2\phi)}{n \cdot \kappa}\right)$, which simplifies to $O\left(\frac{\tau + 2\phi}{\kappa}\right)$. Because τ and ϕ are constants, the relative overheads are $\frac{O(1)}{\kappa}$, as claimed.

It remains to bound the overheads involved in the evaluation of the spguards that contain less than $\Theta(\kappa)$ sequential work. Each such evaluation induces an overhead $O(1)$, because it costs ϕ , plus τ in case of a parallel execution. We show that the number of these evaluations is at most $O(P \log^2 \kappa)$. To that end, we study *parallel small calls*: evaluations of spguards that involve less than $\Theta(\kappa)$ sequential work and whose parallel body gets executed. Such parallel small calls can be nested, but, due to the balance condition (third assumption from Section 2.4), there can be no more than $O(\log \kappa)$ levels of nesting. As we argue in the long version of the paper, the evaluation of an innermost parallel small call (i.e., one that features no nested parallel small call) necessarily multiplies the value of $N_{\max}$ associated with the estimator by a factor α . Thus, for each syntactic spguard and for each processor, the number of innermost parallel small calls on that estimator is $O(\log \kappa)$. Given that a program contains a constant number of syntactic spguards, we deduce that there are at most $O(P \log \kappa)$ innermost parallel small calls during an execution. Then, because each parallel small call contains at least one innermost parallel small call, and because an innermost parallel small call is nested in at most $O(\log \kappa)$ parallel small calls, we deduce that the total number of parallel small calls is at most $O(P \log^2 \kappa)$. Finally, we argue that the number of *sequential small calls* (i.e., evaluations of spguards that involve less than $\Theta(\kappa)$ sequential work and whose sequential body gets executed) is at most a constant factor larger than the number of *parallel small calls*. By summing up the bounds for parallel small calls and sequential ones, we obtain a $O(P \log^2 \kappa)$ bound for the overheads associated with all calls to spguard involving less than $\Theta(\kappa)$ sequential work. This concludes the proof summary.

3 Experimental evaluation

This section presents a study of eight benchmarks from The Problem Based Benchmark Suite (PBBS) [Blelloch et al. 2012]. These benchmarks consist of state-of-the-art solutions to problems on sequences, strings, graphs, and in geometry and graphics. The implementation of the benchmarks consists of a collection C++ codes that use the linguistic extensions of Cilk Plus to realize parallel execution. Parallelism is expressed across the source codes in the idiomatic style of Cilk, namely, nested parallelism, involving, in specific, parallel loops occurring inside parallel loops mixed with calls to recursive, parallel functions. As is common in nested parallelism, the applications generate irregular and, often, highly data-dependent parallelism, thereby making granularity control a particular challenge. Two of the benchmarks

we consider, namely ray-cast and delaunay, represent the largest and most complex codes in PBBS.

The original, authors' codes leverage a variety of manual techniques to control granularity. These techniques involve program instrumentation along with careful tuning to achieve efficiency. For our programs, we ported the original PBBS codes by inserting spguards where necessary, and, for the spguards, we wrote a total of 24 explicit cost functions. It was otherwise possible to use default cost functions, as described in Section 2.6. We summarize the changes:

1. A few divide-and-conquer algorithms involved a manually fixed grain to control granularity. We replaced each of them with a spguard, thereby eliminating the hardware-dependent magic numbers.
2. A number of loops were parallelized by splitting in fixed-size blocks of 2048 items each—a pattern widely used throughout the PBBS sequence library. We replaced all of them with our automatically controlled parallel-for loops.
3. Other loops exploited Cilk parallel for-loops, which essentially split the loop range in $8P$ blocks, where P is the number of cores. We also replaced all of them with our automatically controlled loops.
4. A number of inner loops were forced to sequentialize, even though they could have been parallel. As we learned from private communication with the authors, the purpose was to tame the overheads. We restored parallelism using our automatically-controlled loops.
5. Two loops were forced to always make one spawn per iteration, using a Cilk for-loop annotated with `#pragma grainsize = 1`. Again, we replaced all such loops with our automatically-controlled loops, which properly support nested parallelism.

Overall, automatic granularity control enables replacing careful selection of techniques, such as those listed above, with a single, uniform technique for controlling granularity, able to handle nontrivial nesting of parallelism constructs. Furthermore, the technique automates away the labor-intensive tuning of grain sizes, and it automatically adapts to the hardware—performance is portable.

Yet, these benefits are not completely free. The programmer must take care to ensure that the algorithm in question meets the requirements of our approach, as described in Section 2. Owing to such requirements, we left out eight of the remaining PBBS benchmarks. Fortunately, for each such benchmark, the reason for the incompatibility is straightforward, and in the long version of the paper, we explain the reasons. In spite of such cases, our approach nevertheless remains sound, even in applications where there are some functions that use our approach, and some that use, for example, manual granularity control, or perhaps some other method. The only caveat is that the guarantees associated with our approach do not apply for regions of the program

in which there are calls to manually-controlled functions from inside spguards.

In our approach, there are potential overheads related to the need to infer granularity thresholds online, through a convergence phase. Nevertheless, as established by our analysis (and, as we confirm through our experimental results), the cost of the convergence phase generally accounts for at most a few percent of the parallel run time. Overheads associated with manual granularity control pose their own challenges, but are not backed by guarantees. We observed that, by removing the pragmas (fifth item in the list above), performance would sometimes degrade by up to 5%, and we note that the degradation observed by the PBBS authors may have been more pronounced. The overheads associated with the fourth item in the list above are analyzed in detail by our study of the parallel BFS benchmark, which shows that our automatic method achieves better performance across a wider range of inputs than any manual setting.

Experimental setup. Our primary test harness is an Intel machine with 40 cores. We compiled the code using GCC (version 6.3) using the extensions for Cilk Plus (options `-O2 -march=native -fcilkplus`). Our 40-core machine has four 10-core Intel E7-4870 chips, at 2.4GHz, with 32Kb of L1 and 256Kb L2 cache per core, 30Mb of L3 cache per chip, and 32Gb RAM, and runs Ubuntu Linux kernel v3.13.0-66-generic. For each data point, we report the average running time over 30 runs. The variation in the running times is negligible: overall, we observed only a few cases where the standard deviation is 5%, but it was usually below 3%.

To pick settings for parameters κ and α , we developed a method that involves a one-time-per-machine, automatic tuning step. The starting point is a benchmark program that takes as input an array of 32-bit integers and computes the sum using a parallel reduction. The reduction code uses a spguard to control granularity. To pick κ , we perform a series of runs of the reduction benchmark using a single core, starting with a small setting of $\kappa = 1\mu\text{s}$ and progressively trying larger settings. For these initial runs, we pick $\alpha = 1.3$, a small setting so that the estimator has plenty of slack to converge to reach the target κ . We stop this process upon reaching the first setting of κ for which the program runs slightly above 5% slower than the *sequential elision* of the our reduction benchmark. The sequential elision is (in general) the version of the parallel program in which all fork-join and spguard constructs are erased, leaving only the bare sequential code.

Having picked κ , we then run the same benchmark, but this time using all available cores and trying a number of settings in $1.3 \leq \alpha \leq 5.0$. We then pick α to be the setting that gives the best running time. On our test machine, this process gave us $\kappa = 10.2\mu\text{s}$ and $\alpha = 3.0$.

1	2	3	4	5	6	7	8	9
Application/input	Sequential elision		1-core execution		40-core execution			
	PBBS	Oracle	PBBS	Oracle	PBBS	Oracle	Oracle / PBBS	
	(s)		(relative to elision)		(s)		Idle time	Nb threads
samplesort								
random	21.765	+7.1%	+15.6%	-1.9%	0.834	-7.0%	-7.1%	-18.1%
exponential	15.682	+5.3%	+13.1%	-1.4%	0.626	-4.4%	-4.4%	-15.7%
almost sorted	6.799	+19.3%	+26.4%	-4.3%	0.423	-7.8%	-7.8%	-3.0%
radixsort								
random	3.060	+3.6%	+0.5%	-5.3%	0.235	+8.8%	+9.3%	-11.4%
random pair	4.920	-0.3%	+0.3%	-1.0%	0.469	+4.3%	+3.2%	+128.0%
exponential	3.135	+3.3%	+1.0%	-2.6%	0.245	-1.9%	-3.1%	+23.7%
suffixarray								
dna	19.276	+1.4%	+2.2%	+3.5%	1.494	+3.8%	+1.4%	+28.2%
etext	70.322	+1.1%	+2.0%	+2.4%	4.527	-0.0%	-1.7%	+68.4%
wikisamp	62.112	-1.1%	+2.0%	+3.7%	4.139	+3.2%	+1.4%	+64.1%
convexhull								
kuzmin	7.040	+24.5%	-0.8%	-15.1%	0.517	-0.1%	+6.0%	-81.6%
on circle	130.835	-6.1%	+104.8%	+6.7%	9.871	-31.5%	-31.5%	-23.4%
nearestneighbors								
kuzmin	19.712	+1.8%	+21.7%	+13.2%	1.308	-1.8%	-3.5%	+20.0%
plummer	27.173	-3.4%	+6.2%	+3.4%	2.444	-4.6%	-10.5%	-0.3%
delaunay								
in square	63.967	-0.1%	-0.6%	-0.4%	3.150	-0.6%	-0.8%	+3.2%
kuzmin	69.428	+0.7%	+0.5%	+1.7%	3.802	-4.0%	+0.9%	-21.4%
raycast								
happy	10.698	+6.3%	+3.2%	+0.3%	0.473	+5.5%	+5.5%	-32.7%
xyzrgb	338.920	-1.9%	-2.5%	+1.2%	9.677	+1.5%	+2.0%	-67.3%

Table 1. Benchmark results. Column 3 gives an estimate of the difference in performance between our implementation and the original PBBS version, on a single core, neglecting parallelism-related overheads. Column 4 gives a lower bound on the overheads of the original PBBS code, with figures relative to Column 2. Column 5 gives an estimate of the thread-creation and spguard overheads in our approach, with figures relative to Column 3. In the 40-core section, Column 6 gives PBBS execution time, and Column 7 gives the oracle-guided figure relative to Column 6. Negative figures indicate that our approach is performing better. Columns 8 and 9 give the ratios between our approach and that of PBBS for two measures: total idle time and number of threads created.

3.1 Main PBBS results

For input data sets, we reused much from the original PBBS study [Blleloch et al. 2012], but introduced newly acquired data in a few cases. The inputs to *comparison-* and *radix-sort* consist of sequences of length 10^8 . For *convex-hull*, the inputs consist of 10^8 2-d points. For *nearest neighbors*, the *kuzmin* input consists of 10^8 2-d points and *plummer*, 10^8 3-d points. For *delaunay*, the inputs consist of 10^7 2-d points. For *ray-cast*, we used a two non-synthetic inputs: *happy* consists of the Happy Buddha mesh from the Stanford 3D Scanning Repository, and it consists of 1 million triangles, and *xyz-rgb-manuscript* comes from the same repository and consists of 4 million triangles. Finer details appear in the long version of the paper.

The main performance results appear in Table 1. Columns 2 and 3 represent runs of the sequential elision of the original

PBBS code and of our implementation, respectively. The results in column 3 in particular show the relative performance of our implementation versus the original PBBS one. A -5% value indicates that our version performs 5% faster. Although our code is almost always identical to the original PBBS code, there are some minor structural differences in the function bodies that affect performance slightly. Overall, however, the sequential elisions in both cases perform similarly.

We next evaluate the overheads of thread creation in the original PBBS codes and in our oracle-guided versions. To evaluate for original PBBS, we compare the execution time of the sequential elision of the PBBS code to that of the PBBS code compiled as a Cilk parallel binary. The estimation of overheads in Cilk may be incomplete because the Cilk system in some places detects at runtime that there is only one active worker thread in the system. Nevertheless, the comparison

should give us a lower bound on the parallelism overheads affecting Cilk programs. Column 4 in Table 1 shows that overheads are sometimes low, but sometimes over 15%, and in one case 104%.

To examine overheads of our approach, we compare in a similar fashion the execution time of our sequential elision to that of our parallel-ready binary, running on a single core. Our parallel-ready binary includes the overhead costs of both the Cilk parallel constructs and the estimation-related overheads of the granularity controller. Column 5 in Table 1 shows that the parallel-ready binary is in one case 13% slower, but is otherwise close to the overhead of 5% or smaller that we targeted. There are a number of cases where the parallel-ready binary is slightly faster than the sequential elision. The reason is that the sequential elision switches from recursive, parallel-ready to loop-based sequential code when the complexity function returns a value smaller than 10k, which is sometimes slightly less efficient than the threshold given by the parallel-ready binary.

We finally compare the execution time of original PBBS to that of our binary on 40 cores, using automatic granularity control. Columns 6 and 7 in Table 1 show the results. The results show that in all but two cases, our binary is faster or on par with the PBBS binary. When it is slower, our binary is slower by 8.8%, whereas when it is faster, it is in one case 31.5% and another 7.8% faster. These latter two cases correspond to worst-case inputs in samplesort and convex-hull, where the original PBBS binary is slower because of comparatively poor utilization of the cores.

To gain further insight, we included Columns 7 and 8 in Table 1. Column 8 shows the ratio between idle time (counting periods during which workers are out of work) in our approach and idle time in PBBS. The figures show that the idle time is of the same order of magnitude.² Column 9 shows the ratio between the number of threads created in our approach and the number of threads created in PBBS. Taken together, these last two columns indicate that our approach is able to achieve similar utilization despite having to modulate the creation of threads online.

3.2 Parallel BFS

In addition to the other PBBS programs, we considered the non-deterministic BFS benchmark, named ndBFS in the PBBS paper. We chose ndBFS because it is the fastest among the alternative BFS implementations in PBBS and, to the best of our knowledge, the fastest publicly available Cilk implementation of BFS. To test of the robustness of granularity control, we picked the input graphs to feature a wide range of graph structures. These graphs were used in a different

performance study to test the granularity control of a DFS-like graph-traversal algorithm [Acar et al. 2015]. We picked a representative subset of the graphs, including small-world graphs, such as livejournal, twitter, and wikipedia, and high-diameter graphs, such as europe and cube-grid.

There are two versions of BFS: the *flat* and the *nested* version. In the flat version, the algorithm traverses over the list of neighbors of a vertex in the BFS frontier sequentially, whereas, in the nested version, the algorithm uses a parallel loop, thereby allowing to process in parallel the out-edges of vertices with high out-degree.

PBBS currently uses the flat version. Via personal communication, we learned that the authors sequentialized the inner loop because (1) the graphs they considered did not feature such high-degree vertices, and (2) they observed that performance improved for some of their test graphs when this loop was serialized. Columns 2 and 6 from Table 2 give the execution time for the flat PBBS BFS, and for its nested counterpart, using a parallel `cilk_for` loop over the edges. The figures from these two columns confirm that overheads of parallelization using Cilk-for are significant, sometimes above 50%.

In contrast, our algorithm supports nested parallelism and, even so, delivers significantly better results for nested PBBS, as reported in Column 7 from Table 2 (“nested - ours”). The last column from this figure shows our main result: it compares the performance of the original authors’ flat BFS versus our nested BFS with automatic granularity control. Our version either performs about as well or up to 84% faster.

3.3 Summary

Overall, the results show that our technique enables a simple, uniform method for granularity control, whereas the original PBBS code involves several types of manual intervention. The main PBBS results show our method delivering similar or better performance overall, in spite of having to adaptively adjust grain settings on the fly. The BFS results show our technique delivering consistently better performance in the face of fine-grain, nested, and irregular parallelism. In the long version of the paper, we present a portability study, showing that the technique yields similar results across various test machines with different core counts and architectures.

4 Related work

Controlling the overheads of parallelism has been an important open problem since the early 1980’s, when it was first recognized that practical overheads can undo the benefits of parallelism [Halstead 1984; Mohr et al. 1990]. Since then, researchers have explored two separate approaches.

Granularity Control. Perhaps the oldest technique for granularity control is to use manual, programmer-inserted “cut-off” conditions that switch from a parallel to a sequential

²Because utilization in these benchmarks is generally between 80% and 99%, the total idle time represents less than 20% of the total execution time, thus a 20% change in idle time would affect the execution time by less than 4%.

	Flat				Nested				Ours nested
Graph	PBBS	Ours	Oracle / PBBS		PBBS	Ours	Oracle / PBBS		vs. PBBS flat
	(sec.)		Idle time	Nb threads	(sec.)		Idle time	Nb threads	
livejournal	0.12	+8.2%	−5.5%	−16.8%	0.19	−29.0%	−2.2%	−58.8%	+11.3%
twitter	2.02	−7.3%	+14.0%	−16.0%	2.71	−32.4%	−0.1%	−57.1%	−9.5%
wikipedia	0.11	+5.1%	−3.9%	−28.1%	0.14	−19.0%	−6.4%	−55.3%	+7.4%
europe	4.28	−38.3%	−50.4%	−24.8%	4.28	−37.9%	−51.3%	−27.6%	−37.9%
rand-arity-100	0.12	+6.5%	−6.8%	+9.8%	0.44	−71.3%	−20.4%	−40.8%	+8.0%
rmat27	0.21	+8.4%	−3.1%	−43.4%	0.41	−46.4%	−5.4%	−56.3%	+4.8%
rmat24	0.36	+5.0%	−3.8%	+4.8%	0.37	+3.0%	−4.6%	+8.8%	+5.0%
cube-grid	0.79	+0.0%	−15.9%	+14.4%	0.79	−0.7%	−17.4%	+10.1%	−0.8%
square-grid	4.51	−38.9%	−24.5%	−31.5%	4.47	−38.1%	−25.5%	−32.7%	−38.7%
par-chains-100	67.7	−71.0%	−71.4%	−90.9%	69.2	−72.6%	−67.7%	−79.6%	−72.0%
trunk-first	13.1	−47.4%	−2.1%	−97.4%	13.4	−48.5%	−2.1%	−97.6%	−47.2%
phases-10-d-2	1.40	+8.1%	−8.9%	−0.4%	0.51	+3.5%	−2.6%	+10.5%	−62.0%
phases-50-d-5	0.67	+5.8%	−5.0%	+6.9%	0.66	+4.7%	−3.6%	+16.6%	+4.3%
trees-524k	12.9	+21.0%	−14.7%	−8.2%	1.80	+15.7%	−2.5%	+23.5%	−83.8%

Table 2. Parallel BFS experiment on 40 cores. Depending on the input graph, either Flat PBBS or Nested PBBS is faster. Our nested-parallel algorithm compares favorably to both of them, as reflected by the last column, and the column “Nested/Ours”.

mode of execution. Researchers have addressed the limitations of such manual granularity control by using various forms of automation. Duran et al. propose a method for selecting among three parallelization options: outer-loop only, inner-loop only, or mixed mode, where inner and outer loops may be parallelized and granularity controlled by a heuristic [Duran et al. 2008]. The cut-off technique of Duran et al. makes its decisions online and, like our technique, requires neither tuning nor recompiling. An earlier approach by Huelsbergen, Larus, and Aiken uses list size to determine the grain size [Huelsbergen et al. 1994], however, their technique assumes linear work complexity for every parallel operation. Another approach uses the height and depth of the recursion tree [Pehoushek and Weening 1990; Weening 1989] to predict the execution time and the grain size, but lacks crucial information because depth and height are not a direct measure of execution time. As Iwasaki et al. point out, because techniques, such as those mentioned here, base decisions solely on certain dynamically collected data, such as recursion-tree depth, height, and dynamic load conditions, the granularity controller typically risks decreasing parallelism adversely [Iwasaki and Taura 2016].

Lopez et al. use a similar approach to ours, but in the context of logic programming [Lopez et al. 1996]. Their cost estimators, however, rely on the asymptotic cost annotations alone, but not on prediction of wall-clock running times. Using the asymptotic cost alone is an overly simplistic method, because, on modern processors, execution time depends heavily on effects of caching, pipelining, etc. Our

technique improves on theirs by defining precise assumptions (Section 2) for correct use, backing with theoretical bounds and with a range of challenging benchmarks.

Iwasaki et al. propose a technique for synthesizing static cut-offs for divide-and-conquer functions in Cilk-style programs [Iwasaki and Taura 2016]. The determination of whether or not to cut to serial is made by using cost estimation from a compiler analysis. As such, this technique relies on having accurate static analysis and sophisticated compiler support. Our technique offers an alternative that can, in contrast to theirs, be implemented entirely as a library, using off-the-shelf tools, such as GCC, and can adapt its decisions to changing conditions rather than compile-time estimates. The main cost of which is sometimes having to write cost functions.

Lazy task creation. Lazy task creation, lazy scheduling, and heartbeat scheduling are techniques that aim to reduce the number of tasks created by observing the load in the system and creating tasks only when necessary [Acar et al. 2018; Mohr et al. 1991; Tzannes et al. 2014]. The aim is to reduce the total overhead of thread creation of a parallel program, whereas granularity control aims to reduce the same kind of overhead by instead selectively cutting to serial code.

Lazy task creation has proved to be an indispensable technique for modern parallel programming systems, many of which could not perform well without it. Although it can reduce overheads of parallelism, lazy task creation has some important limitations that prevent it from being the silver bullet. First, there is no way in lazy task creation to account for

the fact that sequential algorithms are usually both asymptotically and practically more work-efficient than their parallel counterparts. At best, lazy task creation can reduce parallelism-related overheads by executing an already parallel algorithm sequentially. Even without parallelism-related overheads, using the stripped-down parallel algorithm (i.e., *sequential elision* [Frigo et al. 1998]) could lead to suboptimal performance, because the problem could have been solved much more efficiently by a faster sequential algorithm. It is thus reasonably common, even when using lazy task creation, to observe a 3-10 fold overhead compared to a piece of code that manually controls granularity of tasks.

Another issue relates to infrastructure requirements. Because it depends on polling on a regular basis for work requests, there is a need for sophisticated support from the compiler [Feeley 1993a]. A key component of such compiler support is the injection of polling checks [Feeley 1993b], which is generally necessary, unless there is support for a suitable inter-core interrupt mechanism à la ADM [Sanchez et al. 2010b], or a suitable interrupt-based alarm mechanism.

We believe, however, that lazy task creation and granularity control, such as our algorithm, are largely complementary: they could easily coexist and provide benefits to each other. Lazy task creation could likely improve performance in cases where a cost function is difficult to define, for example, and our oracle-guided algorithm could likely improve performance by enabling the program to cut in a safe manner to faster serial code. Additional analysis and experimentation will help to better understand the tradeoffs.

5 Conclusion

The problem of managing parallelism-related overheads effectively is an important problem facing parallel programming. The current state-of-the-art in granularity control places the burden of the tuning on the programmer. The tuning process is labor intensive and results in highly engineered, possibly performance-brittle code. In this paper, we show that it is possible to control granularity in a more principled fashion. The key to our approach is an online algorithm for efficiently and accurately estimating the work of parallel computations. We show that our algorithm can be integrated into a state-of-the-art, implicitly parallel language and can deliver executables that compete with and sometimes even beat expertly hand-tuned programs.

Acknowledgments. This research is partially supported by the European Research Council (ERC-2012-StG-308246).

A Artifact Appendix

This artifact contains scripts for building the executables, downloading the input data, executing the benchmark programs, and output tables of results similar to Tables 1 and 2.

A.1 Artifact check-list (meta-information)

- Algorithm: samplesort, radixsort, suffixarray, convexhull, nearestneighbors, delaunay, raycast, bfs
- Program: nix build scripts, C++ code
- Compilation: nix build scripts
- Data set: Problem Based Benchmark Suite
- Run-time environment: Linux
- Hardware: An x86-64 machine with multiple cores
- Output: Results tables comparing running times, core idle times, and number of green-thread creations
- How much disk space required (approximately)?: 300GB
- How much time is needed to prepare workflow (approximately)?: 5-10 hours
- How much time is needed to complete experiments (approximately)?: 4 hours
- Publicly available?: Yes
- Code/data licenses (if publicly available?): Yes
- Workflow frameworks used?: nix package manager

A.2 Description

How delivered Our source code consists of a several of our open-source libraries, all of which are under the MIT license. The libraries are hosted along with documentation on Github. The two main repositories are available at <https://github.com/deepsea-inria/sptl> and <https://github.com/deepsea-inria/pbbs-sptl>.

Hardware dependencies The experiments require a machine with multiple x86-64 cores and well-provisioned RAM. Of our test machines, the one with the minimum amount of RAM has 74GB, but this amount is likely more than is needed.

Software dependencies Our build scripts depend on the nix package manager. <https://nixos.org/nix/download.html>

Data sets The data sets take approximately 300GB on disk.

Although they are taken from the Problem Based Benchmark Suite, the data sets used in our experiments are represented in our custom serialized format. As such, the data sets can be accessed only via either the automatic downloading function of our benchmark scripts or the manual downloading process described in Section A.6.

A.3 Installation

After you have the nix package manager, clone our repository and change to the script folder in the clone.

```
$ git clone https://github.com/deepsea-inria/pbbs-sptl.git
```

Then, change to the folder containing the nix build scripts.

```
$ cd pbbs-sptl/script
```

Because the current directory is where the input data is going to be stored, make sure that there is sufficient space on the file system.

Next, to build the benchmarks, run the following command (there are a number of alternative configuration parameters discussed in Section A.6).

```
$ nix-build
```

If the build succeeds, there will be a symlink named `result` in the current directory.

A.4 Experiment workflow

The first step is to run the auto tuner, which picks (machine-specific) settings for κ and α , following the process described in Section 3.

```
$ ./result/bench/autotune
```

The result of this process is to write to `/var/tmp` a folder containing the settings data.

It is important to complete the auto tuning, because otherwise the default settings are used, and the default settings may lead to poor performance on the benchmarks.

To start the main benchmarks, run the following command.

```
$ ./result/bench/bench.pbench compare
```

The command will automatically download the input data sets, and will likely take several hours. We recommend first trying to complete this process on just one benchmark (defaultly, the script will try to run all to completion). See Section A.6 for instructions on how to run individually. If there is a problem downloading the input data automatically, see Section A.6 for instructions on how to obtain the input data manually.

After the previous command finishes successfully, there should appear in the `_results` folder a number of new text files of the form `results_benchmark.txt` and a PDF named `tables_compare.pdf`. The source for the table can be found in the same folder, in the file named `latex.tex`.

To start the BFS benchmarks, run the following command.

```
$ ./result/bench/bench.pbench bfs
```

This command should generate a text file named `results_bfs.txt` and a PDF named `tables_bfs.pdf`. Like before, the sources for the latex table can be found in the same folder, named `latex.tex`.

A.5 Evaluation and expected result

The trends in `tables_compare.pdf` `tables_bfs.pdf` should bear resemblance to the trends in Tables 1 and 2, respectively. The raw results from benchmark runs should appear in text files in the `_results` folder. These results are human readable, but the more efficient way to interpret them is to look at the generated tables.

A.6 Experiment customization

Running benchmarks individually. The benchmarking script, namely `bench.pbench`, supports running one benchmark at a time. For example, the following command runs just `convexhull`.

```
$ ./result/bench/bench.pbench compare -benchmark
convexhull
```

To run multiple, use a comma-separated list, such as the following.

```
$ ./result/bench/bench.pbench compare -benchmark
convexhull,samplesort
```

Manually downloading the input dataset. First, download the input data from <http://mike-rainey.site/ppopp19>. Then, use the following command to build the benchmarks.

```
$ nix-build --argstr pathToData /path/to/data/folder
```

Building with custom modified sources. The interested experimenter may want to edit the source code and benchmark with their changes. Modifications to three critical repositories used by our benchmarks can be built by the following process. First, from

the folder containing `pbbs-sptl`, clone the original Problem Based Benchmark Suite library code and our library code.

```
$ git clone https://github.com/deepsea-inria/pbbs-
include.git
```

```
$ git clone https://github.com/deepsea-inria/sptl.git
```

Because the `sptl` repository may change over time, we recommend starting with the revision that is used by the default sources. The revision number can be found in the file at

```
pbbs-sptl/script/default-sources.nix
```

Next, build as follows.

```
$ cd pbbs-sptl/script
```

```
$ nix-build --arg sources 'import ./local-sources.
nix'
```

The library `sptl` contains our granularity-control algorithm along with critical data-parallel library functions used across the benchmarks. Documentation for `sptl` can be built by following the instructions in Section A.6.

Building with library documentation. Documentation for `sptl` is available either in markdown format or in easier-to-read PDF and HTML formats. To render PDF and HTML documentation, pass to the build script the `buildDocs` flag as follows.

```
$ cd pbbs-sptl/script
```

```
$ nix-build --arg buildDocs true
```

After completion, the `sptl` documentation should appear along with the markdown source in the folder

```
./result/sptl/share/doc
```

Running with custom κ , α settings The following command line overrides the settings generated by the auto tuner.

```
$ ./result/bench/bench.pbench compare -sptl_kappa
12.0 -sptl_alpha 3.0
```

Picking the number of processors manually The number of cores in the machine is detected by the autotuner, and that number is used by default for benchmarking. To override this number, say, to use 12 hardware threads for benchmarking, use the `-proc` flag, as follows.

```
$ ./result/bench/bench.pbench compare -proc 12
```

Picking the number of benchmarking runs The default number of runs per input is 30, and the result reported for every such batch of runs is the average. To override this number to instead perform, say, five runs, use the `-runs` flag, as follows.

```
$ ./result/bench/bench.pbench compare -runs 5
```

Later, to add to the previous results an additional 25 runs per input, use the `-mode append` flag.

```
$ ./result/bench/bench.pbench compare -runs 25 -mode
append
```

A.7 Methodology

Submission, reviewing and badging methodology:

- <http://cTuning.org/ae/submission-20180713.html>
- <http://cTuning.org/ae/reviewing-20180713.html>
- <https://www.acm.org/publications/policies/artifact-review-badging>

References

- Umut A. Acar, Vitaly Aksenov, Arthur Charguéraud, and Mike Rainey. 2019. Provably and Practically Efficient Granularity Control. (2019). <https://hal.inria.fr/hal-01973285v1> Long version.
- Umut A. Acar, Guy E. Blelloch, and Robert D. Blumofe. 2002. The Data Locality of Work Stealing. *Theory of Computing Systems* 35, 3 (2002), 321–347.
- Umut A. Acar, Arthur Charguéraud, Adrien Guatto, Mike Rainey, and Filip Sieczkowski. 2018. Heartbeat Scheduling: Provable Efficiency for Nested Parallelism. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2018)*. ACM, New York, NY, USA, 769–782. <https://doi.org/10.1145/3192366.3192391>
- Umut A. Acar, Arthur Charguéraud, and Mike Rainey. 2013. Scheduling Parallel Programs by Work Stealing with Private Deques. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '13)*.
- Umut A. Acar, Arthur Charguéraud, and Mike Rainey. 2015. A work-efficient algorithm for parallel unordered depth-first search. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2015, Austin, TX, USA, November 15-20, 2015*. 67:1–67:12.
- Umut A. Acar, Arthur Charguéraud, and Mike Rainey. 2016. Oracle-guided scheduling for controlling granularity in implicitly parallel languages. *Journal of Functional Programming (JFP)* 26 (2016), e23.
- Arvind and K. P. Gostelow. 1978. *The Id Report: An Asynchronous Language and Computing Machine*. Technical Report TR-114. Department of Information and Computer Science, University of California, Irvine.
- Lars Bergstrom, Matthew Fluet, Mike Rainey, John Reppy, and Adam Shaw. 2012. Lazy Tree Splitting. *J. Funct. Program.* 22, 4-5 (Aug. 2012), 382–438.
- Guy E. Blelloch, Jeremy T. Fineman, Phillip B. Gibbons, and Julian Shun. 2012. Internally deterministic parallel algorithms can be fast. In *PPoPP '12*. 181–192.
- Guy E. Blelloch and John Greiner. 1996. A provable time and space efficient implementation of NESL. In *Proceedings of the 1st ACM SIGPLAN International Conference on Functional Programming*. ACM, 213–225.
- Guy E. Blelloch, Jonathan C. Hardwick, Jay Sipelstein, Marco Zagha, and Siddhartha Chatterjee. 1994. Implementation of a Portable Nested Data-Parallel Language. *J. Parallel Distrib. Comput.* 21, 1 (1994), 4–14.
- Robert D. Blumofe and Charles E. Leiserson. 1999. Scheduling multithreaded computations by work stealing. *J. ACM* 46 (Sept. 1999), 720–748. Issue 5.
- Manuel M. T. Chakravarty, Roman Leshchinskiy, Simon L. Peyton Jones, Gabriele Keller, and Simon Marlow. 2007. Data parallel Haskell: a status report. In *Proceedings of the POPL 2007 Workshop on Declarative Aspects of Multicore Programming, DAMP 2007, Nice, France, January 16, 2007*. 10–18.
- Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun, and Vivek Sarkar. 2005. X10: an object-oriented approach to non-uniform cluster computing. In *Proceedings of the 20th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications (OOPSLA '05)*. ACM, 519–538.
- A. Duran, J. Corbalan, and E. Ayguade. 2008. An adaptive cut-off for task parallelism. In *2008 SC - International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–11.
- Marc Feeley. 1992. A Message Passing Implementation of Lazy Task Creation. In *Parallel Symbolic Computing*. 94–107.
- Marc Feeley. 1993a. *An efficient and general implementation of futures on large scale shared-memory multiprocessors*. Ph.D. Dissertation. Brandeis University, Waltham, MA, USA. UMI Order No. GAX93-22348.
- Marc Feeley. 1993b. Polling efficiently on stock hardware. In *Proceedings of the conference on Functional programming languages and computer architecture (FPCA '93)*. 179–187.
- Matthew Fluet, Mike Rainey, John Reppy, and Adam Shaw. 2011. Implicitly threaded parallelism in Manticore. *Journal of Functional Programming* 20, 5-6 (2011), 1–40.
- Matteo Frigo, Charles E. Leiserson, and Keith H. Randall. 1998. The Implementation of the Cilk-5 Multithreaded Language. In *PLDI*. 212–223.
- Adrien Guatto, Sam Westrick, Ram Raghunathan, and Umut A. Acar and Matthew Fluet. 2018. Hierarchical Memory Management for Mutable State. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*. ACM Press.
- Robert H. Halstead. 1985. MULTILISP: a language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems* 7 (1985), 501–538.
- Robert H. Halstead, Jr. 1984. Implementation of Multilisp: Lisp on a Multiprocessor. In *Proceedings of the 1984 ACM Symposium on LISP and functional programming (LFP '84)*. ACM, 9–17.
- Tasuku Hiraishi, Masahiro Yasugi, Seiji Umatani, and Taiichi Yuasa. 2009. Backtracking-based load balancing. In *PPoPP '09*. ACM, 55–64. <https://doi.org/10.1145/1504176.1504187>
- Lorenz Huelsbergen, James R. Larus, and Alexander Aiken. 1994. Using the run-time sizes of data structures to guide parallel-thread creation. In *Proceedings of the 1994 ACM conference on LISP and functional programming (LFP '94)*. 79–90.
- Shams Mahmood Imam and Vivek Sarkar. 2014. Habanero-Java library: a Java 8 framework for multicore programming. In *2014 International Conference on Principles and Practices of Programming on the Java Platform Virtual Machines, Languages and Tools, PPPJ '14*. 75–86.
- Intel. 2011. Intel Threading Building Blocks. (2011). <https://www.threadingbuildingblocks.org/>.
- Shintaro Iwasaki and Kenjiro Taura. 2016. A static cut-off for task parallel programs. In *Proceedings of the 2016 International Conference on Parallel Architectures and Compilation*. ACM, 139–150.
- Suresh Jagannathan, Armand Navabi, KC Sivaramakrishnan, and Lukasz Ziarek. 2010. The Design Rationale for Multi-MLton. In *ML '10: Proceedings of the ACM SIGPLAN Workshop on ML*. ACM.
- Gabriele Keller, Manuel M.T. Chakravarty, Roman Leshchinskiy, Simon Peyton Jones, and Ben Lippmeier. 2010. Regular, shape-polymorphic, parallel arrays in Haskell. In *Proceedings of the 15th ACM SIGPLAN international conference on Functional programming (ICFP '10)*. 261–272.
- Doug Lea. 2000. A Java fork/join framework. In *Proceedings of the ACM 2000 conference on Java Grande (JAVA '00)*. 36–43.
- Daan Leijen, Wolfram Schulte, and Sebastian Burckhardt. 2009. The design of a task parallel library. In *Proceedings of the 24th ACM SIGPLAN conference on Object Oriented Programming Systems Languages and Applications (OOPSLA '09)*. 227–242.
- P. Lopez, M. Hermenegildo, and S. Debray. 1996. A methodology for granularity-based control of parallelism in logic programs. *Journal of Symbolic Computation* 21 (June 1996), 715–734. Issue 4-6. <https://doi.org/10.1006/jSCO.1996.0038>
- E. Mohr, D. A. Kranz, and R. H. Halstead. 1991. Lazy task creation: a technique for increasing the granularity of parallel programs. *IEEE Transactions on Parallel and Distributed Systems* 2, 3 (1991), 264–280.
- Eric Mohr, David A. Kranz, and Robert H. Halstead Jr. 1990. Lazy task creation: a technique for increasing the granularity of parallel programs. In *Conference record of the 1990 ACM Conference on Lisp and Functional Programming*. ACM Press, New York, New York, USA, 185–197.
- OpenMP Architecture Review Board. 2008. OpenMP Application Program Interface. (2008). <https://www.openmp.org/>.
- Joseph Pehoushek and Joseph Weening. 1990. Low-cost process creation and dynamic partitioning in Qlisp. In *Parallel Lisp: Languages and Systems*, Takayasu Ito and Robert Halstead (Eds.). Lecture Notes in Computer Science, Vol. 441. Springer Berlin / Heidelberg, 182–199.
- Ram Raghunathan, Stefan K. Muller, Umut A. Acar, and Guy Blelloch. 2016. Hierarchical Memory Management for Parallel Programs. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming (ICFP 2016)*. ACM, New York, NY, USA, 392–406.

- Mike Rainey. 2010. *Effective Scheduling Techniques for High-Level Parallel Programming Languages*. Ph.D. Dissertation. University of Chicago.
- Daniel Sanchez, Richard M. Yoo, and Christos Kozyrakis. 2010a. Flexible architectural support for fine-grain scheduling. In *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems (ASPLOS '10)*. ACM, New York, NY, USA, 311–322.
- Daniel Sanchez, Richard M. Yoo, and Christos Kozyrakis. 2010b. Flexible architectural support for fine-grain scheduling. In *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems (ASPLOS '10)*. ACM, New York, NY, USA, 311–322. <https://doi.org/10.1145/1736020.1736055>
- Alexandros Tzannes, George C. Caragea, Uzi Vishkin, and Rajeev Barua. 2014. Lazy Scheduling: A Runtime Adaptive Scheduler for Declarative Parallelism. *TOPLAS* 36, 3, Article 10 (Sept. 2014), 51 pages. <https://doi.org/10.1145/2629643>
- Joseph S. Weening. 1989. *Parallel Execution of Lisp Programs*. Ph.D. Dissertation. Stanford University. Computer Science Technical Report STAN-CS-89-1265.
- Reza Zadeh. 2017. Overview, Models of Computation, Brent’s Theorem. (2017). https://stanford.edu/~rezab/dao/notes/lecture01/cme323_lec1.pdf