

# A generic algorithm template for divide-and-conquer in multicore systems

Carlos H. González and Basilio B. Fraguela\*

Depto. de Electrónica e Sistemas, Universidade da Coruña, Spain

e-mail: {cgonzalezv, basilio.fraguela}@udc.es

## Abstract

The divide-and-conquer pattern of parallelism has been well known for years. Its recursive nature makes it a powerful approach to organize parallelism on data structures and problems that are expressed naturally in a recursive way. In fact, recent tools such as Intel Threading Building Blocks (TBB), which has received much attention, go further and make extensive usage of this pattern to parallelize problems that other approaches parallelize following other strategies. In this paper we discuss the limitations to express divide-and-conquer parallelism with the algorithm templates provided by the TBB. Based on our observations, we propose a new algorithm template implemented on top of TBB that improves the programmability of many problems that fit this pattern, while providing a similar performance. This is demonstrated with a comparison both in terms of performance and programmability.

Keywords: productivity, programmability, parallel skeletons, template meta-programming, libraries, patterns

## I. INTRODUCTION

The divide-and-conquer strategy appears in many sequential problems [1]. It is applicable whenever the solution to a problem can be found by dividing it into smaller subproblems, which can be solved separately, and merging somehow the partial results to such subproblems into a global solution for the initial problem. This strategy can be often applied recursively to the subproblems until a base or indivisible one is reached, which is then solved directly. The recursivity of an algorithm sometimes is given by the data structure on which it works, as is the case of algorithms on trees, and very often it is the most natural description of the algorithm. Just to cite a few examples, cache oblivious algorithms [10], many signal-processing algorithms such as discrete Fourier transforms, or the linear algebra algorithms produced by FLAME [3] are usually recursive algorithms that follow a divide-and-conquer strategy. As for parallelism, the independence in the resolution of the subproblems in which a problem has been partitioned leads to concurrency, giving place to the divide-and-conquer pattern of parallelism [15]. A potential problem intrinsically associated with this pattern is that, since often subproblems identified in a given task will be solved by subtasks run in other processors, movements of the data that identify each subproblem are required between processors. While in shared-memory systems this can be achieved easily, it can pose an important problem for the performance in distributed-memory computers, particularly if the amount of data involved is large. In this paper we focus our attention in the recent generations of multicore systems, where the application of divide-and-conquer is more straightforward and has a potential for a better performance.

Although parallelism can always be expressed using low level tools such as POSIX Threads [4], the growth in the size and the complexity of the applications is strongly motivating the usage of higher level approaches that improve the productivity of programmers. This way, fostered by the increase of processors available in current and future systems with the arrival of multicores, extensive research is being made on the best ways to express parallelism. New languages, compiler directives and libraries are the three approaches explored with this aim. The large base of

\* Corresponding author

existing legacy codes, the inherent learning curve and the requirement of compiler support have traditionally made difficult the widespread adoption of new languages focused on parallelism. As for compiler directives, OpenMP [5] is well established in the field of multicore systems. OpenMP is mainly designed to parallelize regular loops, although its scope of application has been extended thanks to the addition of a task-enqueuing mechanism in its latest specification [17]. Compiler directives unfortunately require compiler support and they do not provide as much structure and functionality to applications as libraries of skeletal operations [8]. Skeletons build on parallel design patterns, which provide a clean specification to the flow of execution, parallelism, synchronization and data communications of typical strategies for the parallel resolution of problems. Divide-and-conquer has been in fact identified as one of the basic skeletons of parallel programming [7]. Libraries of parallel skeletons can be thus a very good approach to develop parallel applications thanks to the higher degree of abstraction they provide. A recent example of library that provides parallel skeletons for multicore systems is Intel Threading Building Blocks (TBB) [18], which relies heavily on recursive decomposition of problems and task stealing. Although there have been many other libraries of skeletal operations before TBB, this library has become the most popular and widely adopted. Thus we think it is an interesting contribution to analyze how TBB adapts to the parallelization of typical classes of problems and propose ways to improve it. This way, in this paper (1) we discuss the weaknesses of TBB algorithm templates to parallelize applications that are naturally fit for the divide-and-conquer pattern of parallelism, (2) we propose a new template built on this library to express these problems, and (3) we perform a comparison both in terms of programmability and performance.

The rest of this paper is organized as follows. The next section introduces the TBB, focusing on its ability to express the divide-and-conquer pattern of parallelism. Examples of application to real problems will be discussed in Section III. Our proposal to express this pattern of parallelism is presented in Section IV. An evaluation of our proposal both in terms of performance and programmability is presented in Section V. Related work is discussed in Section VI, followed by our conclusions in Section VII.

## II. THE INTEL TBB LIBRARY

Intel Threading Building Blocks (TBB) [18] is a C++ library developed by Intel for the programming of multithreaded applications. It provides from atomic operations and mutexes to containers specially designed for parallel operation. Still, its main mechanism to express parallelism is the usage of algorithm templates that provide generic parallel algorithms. The most important algorithm templates provided by the TBB are `parallel_for`, `parallel_reduce` and `parallel_scan`, which express element-by-element independent computations on a problem, a parallel reduction, and a parallel prefix or scan operation on a sequence of elements, respectively. These algorithm templates have two compulsory parameters. The first one is a *range* or iteration space that defines the limits of a loop to parallelize, or more in general, a problem that can be recursively subdivided into smaller subproblems that can be solved in parallel. The second one, which we will call *body*, provides the computation to perform on the range. The requirements of the classes of these two objects are now discussed briefly.

The ranges used in the algorithm templates provided by the TBB must model the Range concept, which represents a recursively divisible set of values. Basically, the class must provide

- a copy constructor
- an `empty` method to indicate when a range is empty,
- an `is_divisible` method to inform whether the range can be partitioned into two subranges whose processing in parallel is more efficient than the sequential processing of the whole range,
- and finally a special splitting constructor that splits an input range `r` in two. By convention the constructor should construct the second part of the range, and update `r` (which is an input by reference) to be the first

half. Both halves should be as similar as possible in size in order to attain the best performance.

TBB algorithm templates use these methods to partition recursively the initial range provided into smaller subranges that are processed in parallel. This process, which is completely transparent to the user, seeks to generate enough number of tasks of an adequate size to parallelize optimally the computation on the initial range. Thus, as we see, TBB makes extensive usage of a divide-and-conquer approach to achieve parallelism with its templates. This recursive decomposition is complemented by a task-stealing scheduling which allows to balance the load among the existing threads, generating and moving subtasks among them as needed.

The class that provides the body or computation on the range has different requirements depending on the algorithm template. This way, `parallel_for` only requires that the class has a copy constructor and overloads the `operator()` method on the range class used. The parallel computation is performed in this method. `parallel_reduce` requires additionally a splitting constructor and a `join` method. The splitting constructor is used to build copies of the body object for the different threads that participate in the parallel reduction. The `join` method has as input a `rhs` body that contains the reduction of a subrange just to the right (i.e following) the subrange reduced in the current body. The method must update the object on which it is invoked to represent the accumulated result for its reduction and the one in `rhs`, that is, `left.join(right)` should update `left` to be the result of `left` reduced with `right`. Notice that the reduction operation should be associative, but it need not be commutative. It is important that a new body is created only when a range is split, but the converse is not true. This means that a range can be recursively subdivided in several smaller subranges which are all reduced by the same body in its `operator()` method. When this happens, the subranges are always evaluated by the body in left to right order, so that non commutative operations are not endangered.

We will not detail the requirements of the body of the `parallel_scan` algorithm template because it was not of interest for the resolution of any of the problems considered in this paper. This way in what follows we will focus on the usage of `parallel_for` and `parallel_reduce` to express the divide-and-conquer pattern of parallelism. As we have seen, both algorithm templates subdivide recursively their input problem in smaller ones that can be processed simultaneously. The latter additionally gathers the partial results corresponding to several subproblems and combines them to compute the solution for the initial larger problem they conform.

Finally, TBB algorithm templates allow a third optional parameter called the *partitioner*. This object indicates the policy to follow to generate new parallel tasks. When it is not provided, it defaults to the `simple_partitioner`, which recursively splits the ranges giving place to new subtasks until their `is_divisible` method returns `false`. Thus in this case the programmer is fully in control of the generation of parallel tasks. The `auto_partitioner` lets the TBB library decide whether the ranges must be split to balance load in an automated way. Thus ranges could not be split even if they were divisible because their division is not needed to balance load. The third partitioner is `affinity_partitioner`, which is designed for algorithms that are performed several times on the same data and this data can fit in the caches. It tries to assign the same iterations of loops to the same threads that run them in a past execution of the loop.

### III. DIVIDE-AND-CONQUER WITH THE TBB

This Section analyzes the programmability of the divide-and-conquer pattern of parallelism using the TBB algorithm templates through a series of examples of increasing complexity. This analysis motivates and leads to the design of the alternative that will be presented in the next Section.

```

1 struct FibRange {
2   int n_;
3
4   FibRange(int n)
5   : n_(n) { }
6
7   FibRange(FibRange& other, split)
8   : n_(r.n_ - 2)
9   { other.n_ = other.n_ - 1; }
10
11  bool is_divisible() const { return (n_ > 1); }
12
13  bool empty() const { return n_ < 0; };
14 };

```

↔

```

16 struct Fib {
17   int fsum_;
18
19   Fib()
20   : fsum_(0) { }
21
22   Fib(Fib& other, split)
23   : fsum_(0) { }
24
25   void operator() (FibRange& range) { fsum_ += fib(range.n_); }
26
27   int fib(int n) {
28     if (n < 2) return n;
29     else return fib(n - 1) + fib(n - 2);
30   }
31
32   void join(Fib& rhs) { fsum_ += rhs.fsum_; };
33 };
34 ...
35 Fib f();
36 parallel_reduce(FibRange(n), f, auto_partitioner());
37 int result = f.fsum_;

```

Figure 1. Computation of the  $n$ -th Fibonacci number using TBB's `parallel_reduce`

### A. Fibonacci numbers

The simplest program we will consider is the recursive computation of the  $n$ th Fibonacci number. While this is an inefficient method to compute this value, our interest at this point is not on performance, but on the expressiveness of the library, and this problem is ideal because of its simplicity. The sequential version is

```

int fib(int n) {
  if (n < 2) return n;
  else return fib(n - 1) + fib(n - 2);
}

```

which clearly shows all the basic elements of a divide and conquer algorithm:

- the identification of a base case (when  $n < 2$ )
- the resolution of the base case (simply return  $n$ )
- the partition in several subproblems otherwise ( $\text{fib}(n - 1)$  and  $\text{fib}(n - 2)$ )
- the combination of the results of the subproblems (here simply adding their outcomes)

The simplest TBB parallel implementation of this algorithm, shown in Figure 1, is based on `parallel_reduce` and it indeed follows a recursive divide-and-conquer approach. The range object required by this template is provided by the `FibRange` class, and it encapsulates which is the Fibonacci number to compute in `n_`. The body object belongs to the `Fib` class and performs the actual computation. These classes implement the methods required by the template, which were explained in Section II. For example the splitting constructors are the constructors that have a second dummy argument of type `split`. The splitting constructor for `FibRange` in lines 7-9 shows how the computation of the  $n$ -th Fibonacci number is split in the computation of the  $n - 1$  and  $n - 2$ th numbers. Concretely, line 8 fills in the new range built to represent the  $n - 2$ th number, while the `FibRange` which is being split, called `other` in the code, represents now the  $n - 1$ th number (line 9). Notice also that the `operator()` of the body must support any value of the input range, and not just a not divisible one (0 or 1). The reason is that the `auto_partitioner` is being used (line 36) to avoid generating too many tasks. The default `simple_partitioner` would have generated a new task for every step of the recursion, which would have been very inefficient.

The `Fib` class must have a state for three reasons. First, the same body can be applied to several ranges, so it must accumulate the results of their reductions. Second, bodies must also accumulate the results of the reductions of other bodies through the calls to their `join` method. Third, TBB algorithm templates have no return type, thus

```

1 struct TreeAddRange {
2   tree_t * r1_, *r2_;
3   tree_t * parent_;
4
5   TreeAddRange(tree_t *root)
6   : r1_(root), r2_(0), parent_(0) { }
7
8   TreeAddRange(TreeAddRange& other, split) {
9     if(other.r2_ == 0) { //other only has a node
10      r1_ = other.r1_->child[1];
11      r2_ = other.r1_->child[2];
12      parent_ = other.r1_;
13      other.r1_ = other.r1_->child[0];
14    } else { //other has two nodes
15      parent_ = 0;
16      r1_ = other.r2_;
17      r2_ = 0;
18      other.r2_ = 0;
19    }
20  }
21
22  bool empty() const { return r1_ == 0; }
23
24  bool is_divisible() const { return !empty(); }
25 };

```

↔

```

27 struct TreeAddReduce {
28   int sum_;
29
30   TreeAddReduce()
31   : sum_(0) { }
32
33   TreeAddReduce(TreeAddReduce& other, split)
34   : sum_(0) { }
35
36   void operator()(TreeAddRange &range) {
37     sum_ += TreeAdd(range.r1_);
38     if(range.r2_ != 0)
39       sum_ += TreeAdd(range.r2_);
40     if (range.parent_)
41       sum_ += range.parent_>val;
42   }
43
44   void join(TreeAddReduce& rhs) {
45     sum_ += rhs.sum_;
46   }
47 };
48 ...
49 TreeAddReduce tar;
50 parallel_reduce(TreeAddRange(root), tar, auto_partitioner());
51 int r = tar.sum_;

```

Figure 2. Reduction on a 3-ary tree using TBB’s `parallel_reduce`

the body objects must store the results of the reductions. This gives place to the invocation we see in lines 35-37 of the listing. The topmost `Fib` object must be created before the usage of `parallel_reduce` so that when it finishes the result can be retrieved from it.

Altogether, even when the problem suits well the TBB algorithm templates, we have gone from 4 source lines of code (SLOC) in the sequential version to 26 (empty lines and comments are not counted) in the parallel one.

## B. Tree reduction

TBB ranges can only be partitioned in two subranges in each subdivision, while there are many problems in which it is more natural to divide a task in more than two subtasks. For example, the natural representation of a subproblem in a parallel operation on a tree is a range object that stores a single node. When this range is processed by the `operator()` method of the body of the algorithm template, the node and its children are processed. This way, in a parallel operation on an 3-ary tree, the task/range corresponding to the parallel operation on a node would naturally be subdivided in 3 subtasks/ranges, one per direct child. The TBB restriction to just two subranges in each partition forces the programmer to build a more complex representation of the problem so that there are range objects that represent a single child node, while others keep two children nodes. As a result, the construction and splitting conditions for both kinds of ranges will be different, implying a more complicated implementation of the methods of the range. Moreover, the `operator()` method of the body will have to be written to deal correctly with both kinds of ranges.

These problems are exemplified in Figure 2 with the TBB implementation of a reduction on a 3-ary tree. The initial range stores the root of the tree in `r1_`, while `r2_` is set to 0. The splitting constructor operates in a different way depending on whether the range `other` to split has a single node or two (line 9). If it has a single node, the new range takes its two last children and stores that its parent is `other.r1_`. The input range `other` is then updated to store its first child. When `other` has two nodes, the new range takes the second one and zeroes it from `other`. The `operator()` of the body for the reduction has to take into account whether the input range has one or two nodes, and also whether a parent node is carried.

This example also points out another two problems of this approach. Although a given task that can be subdivided in  $N > 2$  subtasks can always be subdivided only in two, those two subproblems may have necessarily a very different granularity. In this example, one of the two children ranges of a 3-ary node is exactly twice larger than the other. This can lead to a poor load balancing, since the library recommends the subdivisions to be as even as possible. This problem can be alleviated by further subdividing the ranges and relying on the work-stealing scheduler of the TBB, which can move tasks from loaded processors to idle ones. Still, the TBB does not provide a mechanism to specify which ranges should be subdivided with larger priority, but just a boolean flag that indicates whether a range can be further subdivided or not. Moreover, when automatic partitioning is used, the library may choose not to split a range even if it is divisible. For these reasons, allowing to subdivide in  $N$  subranges at once improves both the programmability and the potential performance of divide-and-conquer problems.

The last problem is the difficulty to handle correctly pieces of the description of the problem which are not a natural part of the representation of the children subproblems, but which are required in the reduction stage. In this code the problem is reflected by the clumsy treatment of the inner nodes of the tree, which must be stored in the `parent_` field of the ranges taking care that none is either lost or stored in several ranges. Additionally, the fact that some ranges carry an inner node in this field while others do not complicates the `operator()` of the body, as we have said.

### C. Traveling salesman problem

An important problem that the TBB template algorithms pose for many classes of divide-and-conquer algorithms is that they require the reduction operations on the ranges to be associative. This complicates the implementation of the algorithms in which the solution to a given problem at any level of decomposition requires merging exactly the solutions to its children subproblems. An algorithm of this kind is the recursive partitioning algorithm for the traveling salesman developed in [13], an implementation of which is the *tsp* Olden benchmark [19]. The program first builds a binary space partitioning tree with a city in each node. Then the solution is built traversing the tree with the function

```
Tree tsp(Tree t, int sz) {
    if (t->sz <= sz) return conquer(t);
    Tree leftval = tsp(t->left, sz);
    Tree rightval = tsp(t->right, sz);
    return merge(leftval, rightval, t);
}
```

As we see, it follows a typical divide-and-conquer strategy. The base case is found when the problem is smaller than a given size `sz` and it is solved with the function `conquer`. Otherwise the two children can be processed in parallel applying `tsp` recursively. The solution is obtained joining their solutions (tours) with the `merge` function. This requires inserting the parent node in the solution, as it corresponds to the city located in the cut in the plane that gave place to the two children.

This structure fits well the `parallel_reduce` template in many aspects. The range and body classes used for the parallelization with this algorithm template are shown in Figure 3. The range contains a node, and splitting it returns the two children subtrees. The `is_divisible` method checks whether the subtree is smaller than `sz`, when the recursion stops. The `operator()` of the body applies the original `tsp` function on the node taken from the range.

The problems arise when the application of the `merge` function to join the partial solutions is considered. First, a stack must be added to the range for two reasons. One is to identify when two ranges are children of the same

```

1 struct TSPRange {
2     static int sz_;
3     stack<Tree> ancestry_;
4     Tree t_;
5
6     TSPRange(Tree t, int sz)
7     : t_(t)
8     { sz_ = sz; }
9
10    TSPRange(TSPRange& other, tbb::split)
11    : t_(other.t_>right), ancestry_(other.ancestry_)
12    {
13        ancestry_.push(other.t_);
14        other.ancestry_.push(other.t_);
15        other.t_ = other.t_>left;
16    }
17
18    bool empty() const { return t_ == 0; }
19
20    bool is_divisible() const { return (t_>sz > sz_); }
21
22    bool mergeable(const TSPRange& rhs) const {
23        return !ancestry_.empty() && !rhs.ancestry_.empty() &&
24            (ancestry_.top() == rhs.ancestry_.top());
25    }
26 };

```

↔

```

28 struct TSPBody {
29     list<TSPRange> lresults_;
30
31     TSPBody() { }
32
33     TSPBody(TSPBody& other, tbb::split) { }
34
35     void operator() (TSPRange &range) {
36         range.t_ = tsp(range.t_, range.sz_);
37         mergeTSPRange(range);
38     }
39
40     void mergeTSPRange(TSPRange &range) {
41         while (!lresults_.empty() && lresults_.back().mergeable(range)) {
42             range.t_ = merge(lresults_.back().t_, range.t_, range.ancestry_.top());
43             range.ancestry_.pop();
44             lresults_.pop_back();
45         }
46         lresults_.push_back(range);
47     }
48
49     void join(TSPBody& rhs) {
50         list<TSPRange>::iterator itend = rhs.lresults_.end();
51         for(list<TSPRange>::iterator it = rhs.lresults_.begin(); it != itend; ++
52             it) {
53             mergeTSPRange(*it);
54         }
55     };
56     ...
57     parallel_reduce(TSPRange(root, sz), TSPBody(), auto_partitioner());

```

Figure 3. Range and body for the Olden tsp parallelization using TBB's `parallel_reduce`

parent and can thus be merged. This is expressed by function `mergeable` (lines 22-25). The other reason is that this parent is actually required by `merge`.

The reductions on the results of the computations of the ranges take place in two places. First, according to the TBB specification, a body operator() can be applied to several consecutive ranges in left to right order, and must reduce their results. This way, when `tsp` is applied to the node in the input range (line 36), the result is stored again in this range and an attempt to merge it with the results of ranges previously processed is done in method `mergeTSPRange` (lines 40-47). The body keeps a list `lresults_` of ranges already processed with their solution. The method repetitively checks whether the rightmost range in the list can be merged with the input range. In this case, `merge` reduces them into the input range, and the range just merged is removed from the list. In the end, the input range is added at the right end of the list. Reductions also take place when different bodies are accumulated in a single one through their `join` method. Namely, `left.join(right)` accumulates in the left body its results with those of the right body received as argument. This can be achieved applying `mergeTSPRange` to the ranges in the list of results of the rhs body from left to right (lines 49-54).

#### IV. AN ALGORITHM TEMPLATE FOR DIVIDE-AND-CONQUER PROBLEMS

The preceding Section has illustrated the problems that programmers face to parallelize many divide-and-conquer problems with the TBB. The restriction to binary subdivisions or the lack of associativity lead, not surprisingly, to an increase in the complexity of the implementation. But even problems that seem to fit well the TBB paradigm such as the recursive computation of the Fibonacci numbers have a large parallelization overhead, as several kinds of constructors are required, reductions on the outcomes of the processing of subproblems can take place in different places, bodies must keep a state to perform those reductions, etc.

The basic components of a divide-and-conquer algorithm are the identification of the base case, its resolution, the partition in subproblems of a non base problem, and the combination of the results of the subproblems. Thus

```

template<typename T, int N>
struct Info : Arity<N> {
    bool is_base(const T& t) const; //is t the base case of the recursion?
    int num_children(const T& t) const; //number of subproblems of t
    T child(int i, const T& t) const; //get i-th subproblem of t
};

template<typename T, typename S>
struct Body : EmptyBody<T, S> {
    void pre(T& t); //preprocessing of t before partition
    S base(T& t); //solve base case
    S post(T& t, S *r); //combine children solutions
};

```

Figure 4. Templates that provide the pseudo-signatures for the info and body objects used by `parallel_recursion`

we should try to provide a way to express these problems using just one method for each one of these components. In order to increase the flexibility, the partition of a non-base problem could be split in two subtasks: calculating the number of children, so that it is not fixed, and building these children. These tasks could be performed in a method with two outputs, but we feel it is cleaner to use two separate methods for them.

The subtasks identified in the implementation of a divide-and-conquer algorithm can be grouped in two sets, giving place to two classes. The decision on whether a problem is the base case of the recurrence, the calculation of the number of subproblems of non base problems and the splitting of a problem only depend on the input problem data. These methods refer to and manipulate the structure of the input problem. They conform thus an object with a role similar to the range in the standard TBB algorithm templates. We will call this object the *info* object because it provides information on the structure of the problem. Besides, contrary to the range objects, we choose not to encapsulate the problem data inside the info object. This reduces the programmer burden by avoiding the need to write a constructor for this object for most problems.

The processing of the base case and the combination of the solutions of the subproblems of a given problem are responsibility of a second object analogous to the body of the TBB algorithm templates, thus we will call it also body. Many divide-and-conquer algorithms process an input problem of type T to get a solution of type S, so the body must support the data types for both concepts, although of course S and T could be the same. We have found that in some cases it is useful to perform some processing on the input data before checking its divisibility and the corresponding base case computation or recursion. Thus the body of our algorithm template requires a method `pre`, which can be empty, which is applied to the input problem data before any check on it is performed. As for the method that combines the solutions of the subproblems, which we will call `post`, its inputs will be an object of type T, defining the problem at a point of the recursion, and a pointer to a vector with the solutions to its subproblems, so that a variable number of children subproblems is easily supported. The reason for requiring the input problem is that, as we have seen in Sections III-B and III-C, in many cases it has data which are not found in any of its children and which are required to compute the solution.

Figure 4 shows templates that describe the info and body objects required by the algorithm template we propose. The info class must be derived from class `Arity < N >`, where N is either the number of children of each non base subproblem, when this value is a constant, or the identifier `UNKNOWN` if there is not a fixed number of subproblems in the partitions. This class provides a method `num_children` if N is a constant. As for the body, it can be optionally derived from the class `EmptyBody < T, S >`, which provides shell (empty) methods for all the methods a body requires. Thus inheriting from it can avoid writing unneeded methods.

Figure 5 shows the pseudocode for the operation of the algorithm template we propose, which is called `parallel_recursion` for similarity with the names of the standard TBB algorithm templates. Its arguments are the representation of the input problem, the info object, the body object, and optionally a partitioner that defines the policy to spawn parallel subtasks. The figure illustrates the usage of all the methods in the info and body classes, I and B in the figure, respectively. Contrary to the TBB algorithm templates, ours returns a value which has type S, the type of the solution. An specialization of the template allows a return type `void`.

```

1 template<typename T, typename S, typename I, typename B, typename P>
2 S parallel_recursion(T& t, I& i, B& b, P& p) {
3     b.pre(t);
4     if(i.is_base(t)) return b.base(t);
5     else {
6         int n = i.num_children(t);
7         S result[n];
8         if(p.do_parallel(i, t))
9             parallel_for(int j = 0; j < n ; j++)
10                result[j] = parallel_recursion(i.child(j, t), i, b, p)
11     else
12         for(int j = 0; j < n ; j++)
13             result[j] = parallel_recursion(i.child(j, t), i, b, p);
14     }
15     return b.post(t, result);
16 }
17 }

```

Figure 5. Pseudocode of the `parallel_recursion` algorithm template

```

1 struct FibInfo : public Arity<2> {
2     bool is_base(const int i) const { return i <= 1; }
3
4     int child(const int i, const int c) const { return c - i - 1; }
5 };
7 struct Fib: public EmptyBody<int, int> {
8     int base(int i) { return i; }
9
10    int post(int i, int* r) { return r[0] + r[1]; }
11 };
12 ...
13 int result = parallel_recursion<int>(n, FibInfo(), Fib(), auto_partitioner());

```

Figure 6. Computation of the  $n$ -th Fibonacci number using `parallel_recursion`

From the pseudocode we see that `Info::is_base` is not the exact opposite of the `is_divisible` method of the TBB ranges. TBB uses `is_divisible` to express divisibility of the range, but also whether it is more efficient to split the range and process the subranges in parallel than to process the range sequentially. Even if the user writes `is_divisible` to return true for all non base cases, the library can ignore it and stop partitioning even if it indicates divisibility if the `auto_partitioner` is used. For these reasons, the `operator()` of a standard body should be able to process both base and non base instances of the range. This makes it different from the `Body::base` method in Figure 4, which processes the problem if and only if `Info::is_base` is true, as line 4 in Figure 5 shows.

The decision on whether the processing of the children subproblems is made sequentially or in parallel is up to the partitioner in `parallel_recursion` (lines 8-14 in Figure 5). The behavior of the partitioners is as follows. The `simple_partitioner` generates a parallel subtask for each child generated in every level of the recursion, very much as it does in the standard TBB templates after every split. This is the default partitioner. The `auto_partitioner` works slightly different from the standard templates. In them this partitioner can stop splitting the range, even if `is_divisible` is true, in order to balance optimally the load. In `parallel_recursion` this partitioner also seeks to balance load automatically. Nevertheless, it does not stop the recursion in the subdivision of the problem, but just the parallelization in the processing of the subtasks. This way the problem is split in multiple subproblems always that `Info::is_base` is false. Finally, we have added a new partitioner called `custom_partitioner` which takes its decision on parallelization based on an optional `Info::do_parallel(const T& t)` method. If this method returns true, the children of `t` are processed in parallel, otherwise they are processed sequentially.

Let us now review the implementation of the examples discussed in Section III using this proposal.

```

1 struct TreeAddInfo : public Arity<3> {
2   bool is_base(const tree_t *t) const { return t == 0; }
3
4   tree_t *child(int i, const tree_t *t) const { return t->child[i]; }
5 };
7 struct TreeAddBody : public EmptyBody<tree_t *, int> {
8   int base(tree_t * t) { return 0; }
9
10  int post(tree_t * t, int *r) { return r[0] + r[1] + r[2] + t->val; }
11 };
12 ...
13 int r = parallel_recursion<int>(root, TreeAddInfo(), TreeAddBody(),
                                auto_partitioner());

```

Figure 7. Reduction on a 3-ary tree using `parallel_recursion`

```

1 struct TSPInfo: public Arity<2> {
2   static int sz_;
3
4   TSPInfo(int sz )
5   { sz_ = sz; }
6
7   bool is_base(const Tree t) const { return (t->sz <= sz_); }
8
9   Tree child(int i, const Tree t) const { return (i == 0) ? t->left
10  : t->right; }
10 };
12 struct TSPBody : public EmptyBody<Tree, Tree> {
13   Tree base(Tree t) { return conquer(t); }
14
15   Tree post(Tree t, Tree* results) { return merge(results[0], results[1], t); }
16 };
17 ...
18 parallel_recursion<Tree>(root, TSPInfo(sz), TSPBody(), auto_partitioner());

```

Figure 8. Olden tsp parallelization using `parallel_recursion`

### A. Examples of usage

Figure 6 shows the code to compute the  $n$ -th Fibonacci number using `parallel_recursion`. Compared to the 26 SLOC of the implementation based on `parallel_reduce` in Figure 1, this implementation only has 9. This code has the virtue that it not only parallelizes the computation, it even makes unnecessary the original sequential `fib` function thanks to the power of `parallel_recursion` to fully express problems that are solved recursively.

The addition of the values in the nodes of a 3-ary tree, which required 41 SLOC in Figure 2, is expressed using 9 SLOC with `parallel_recursion` in Figure 7. In fact, the version in Figure 2 is a bit longer because it uses the sequential function `TreeAdd`, not shown, to perform the reduction of a subtree in the `operator()` of the body. This function is not needed by the implementation based on `parallel_recursion`, which can perform the reduction just using the template.

Our last example, the traveling salesman problem implemented in the *tsp* Olden benchmark, is parallelized with `parallel_recursion` in Figure 8. The facts that the `post` method that combines the solutions obtained in each level of the recursion is guaranteed to be applied to the solutions of the children subproblems generated by a given problem and that this parent problem is also an input to the method simplify extraordinarily the implementation. Concretely, the code goes from 45 SLOC using `parallel_reduce` in Figure 3 to 12 using `parallel_recursion`.

## V. EVALUATION

We now compare the implementation of several divide-and-conquer algorithms using `parallel_recursion`, the algorithm templates provided in the TBB and OpenMP both in terms of programmability and performance. OpenMP is not directly comparable to the skeleton libraries, as it relies on compiler support. It has been included in this study as a baseline that approaches the minimum overhead in the parallelization of applications for multicores, since the insertion of compiler directives in a program usually requires less restructuring than the definition of the classes that object-oriented skeletons use. This way the comparison of standard TBB and the `parallel_recursion` skeletons with respect to the OpenMP version helps measure the relative effort of parallelization that both kinds of skeletons imply.

The algorithms used in this evaluation are the recursive computation of the  $n$ -th Fibonacci number discussed in Section III-A (`fib`), the merge of two sorted sequences of integers into an output sorted sequence (`merge`), the

Name	Description	Arity	Assoc	SLOC	Effort	V
fib	recursive computation of 43rd Fibonacci number	2	Yes	37	31707	5
merge	merge two sorted sequences of 100 million integers each	2	-	62	143004	6
qsort	quicksort of 10 million integers	2	-	71	119908	11
nqueens	N Queens solution count in $14 \times 14$ board	var	Yes	82	192727	17
treeadd	add values in binary tree with 24 levels	2	Yes	92	179387	9
bisort	sort balanced binary tree with 22 levels	2	No	227	822032	20
health	2000 simulation steps in 4-ary tree with 6 levels	4	No	346	1945582	34
tsp	traveling salesman problem on binary tree with 23 levels	2	No	370	2065129	40

Table I  
BENCHMARKS USED

sorting of a vector of integers using quicksort (qsort), the computation of the number of solutions to the N Queens problem (nqueens) and four tree-based Olden benchmarks [19]. The first Olden benchmark is treeadd, the addition of the values in the nodes of a binary tree, which is very similar to the example in Section III-B, except that the tree has two children per node, which makes it easier to implement using TBB’s `parallel_reduce`. The sorting of a balanced binary tree (bisort), a simulation of a hierarchical health system (health), and the traveling salesman problem (tsp) discussed in Section III-C complete the list.

Table I provides the problem sizes, the number of subproblems in which each problem can be divided (arity) and whether the combination or reduction of the results of the subproblems is associative or not, or even not needed. It also shows the value for a baseline version of the algorithms parallelized with OpenMP of the metrics that will be used in Section V-A to evaluate the programmability. We can see that all the algorithms but nqueens and health are naturally expressed splitting each problem in two, which fits the TBB algorithms templates. The former tries all the possible locations of queens in the  $i$ -th row of the board that do not conflict with the queens already placed in the top  $i - 1$  rows. Each possible location gives place to a child problem which proceeds to examine the placements in the next row. This way the number of children problems varies from 0 to the board size, depending on the number of row and the locations of the queens already placed. As for health, each node in its tree has four subnodes that can be processed in parallel, thus this is the natural number of subproblems. The subnodes are located in a vector. This benefits the TBB algorithm templates, as this enables using as range a `blocked_range`, which is a built-in TBB class that defines a one-dimensional iteration space, ideal to parallelize operations on vectors.

#### A. Programmability

The impact of the use of an approach on the ease of programming is not easy to measure. In this section three quantitative metrics are used for this purpose: the SLOC (source lines of code excluding comments and empty lines), the programming effort [11], and the cyclomatic number [16]. The SLOC is more dependent on the user programming style than the other two metrics. The programming effort is a function of the number of unique operands, unique operators, total operands and total operators found in a program. The operands correspond to the constants and identifiers, while symbols or combinations of symbols that affect the value or ordering of operands constitute the operators. According to [11] the programming effort metric calculated from these values is approximately proportional to the programming effort required to implement an algorithm. Finally, the cyclomatic number [16] is  $V = P + 1$ , where  $P$  is the number of decision points or predicates in a program. The smaller  $V$ , the less complex the program is.

Figure 9 shows the SLOC, programming effort and cyclomatic number increase over an OpenMP baseline version for each code when using a suitable TBB algorithm template (TBB) or `parallel_recursion` (pr). The statistics were collected automatically on each whole application globally. Had we tried to isolate manually the portions

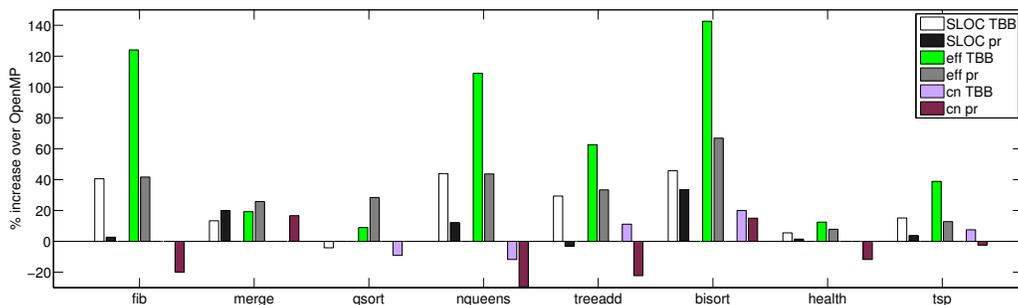


Figure 9. Productivity statistics with respect to the OpenMP baseline version of TBB based (TBB) and `parallel_recursion` based (pr) implementations. SLOC stands for source lines of code, eff for the programming effort and cn for the cyclomatic number.

specifically related to the parallelization, the advantage of `parallel_recursion` over TBB would have often grown to the levels seen in the examples discussed in the preceding sections. We did not do this because sometimes it may not be clear whether some portions of code must be counted as part of the parallelization effort or not, so we measure the whole program as a neutral approach.

The (mostly) positive values indicate that, as expected, OpenMP has the smallest programming overhead, at least when counted with SLOCs or programming effort, which is why it was chosen as baseline. Nevertheless, `parallel_recursion` is the global winner for the cyclomatic number. The reason is that many of the conditionals and loops (they involve conditions to detect their termination) found in divide-and-conquer algorithms are subsumed in the `parallel_recursion` skeleton, while the other approaches leave them exposed in the programmer code more often. `parallel_recursion` requires fewer SLOC, effort and conditionals than the TBB algorithm templates in all the codes but merge and qsort. According to the programming effort indicator, programs parallelized with the TBB templates require 64.6% more effort than OpenMP, while those based on `parallel_recursion` require on average 33.3% more effort than OpenMP. This is a reduction of nearly 50% in relative terms. Interestingly, the situation is the opposite for merge and qsort, in which the average effort overhead over the OpenMP version is 13.4% for the codes that use `parallel_for` and 30.1% for the `parallel_recursion` codes. These are the only benchmarks in which there is no need to combine the result of the solution of the problems: they only require the division in subproblems that can be solved in parallel. They are also the two benchmarks purely based on arrays, where the concept of Range around which the TBB algorithm templates are designed fits better. Thus when these conditions hold, we may prefer to try the standard TBB skeletons.

## B. Performance

The performance of these approaches is compared now using the Intel icpc compiler V 11.0 with optimization level O3 in two platforms. One is a server with 4 Intel Xeon hexa-core 2.40GHz E7450 CPUs, whose results are labeled with X. The other is an HP Integrity rx7640 server with eight dual-core 1.6 GHz Itanium Montvale processors, whose results are labeled with I. Figures 10 to 17 show the performance of the three implementations of the benchmarks on both systems when the automatic partitioning is used in the standard TBB and `parallel_recursion` based approaches. Fib and nqueens use little memory and thus scale well in both systems. The scaling of the other benchmarks is affected by the lack of memory bandwidth as the number of cores used increases, particularly in our Xeon-based system, whose memory bandwidth is up to 5 times smaller than that of the rx7640 server when 16 cores are used. This results in small to null performance improvements when we go from 8 to 16 cores in this system. Benchmark health is also affected by very frequent memory allocations with `malloc` that become a source of contention for the threads due to the associated lock.

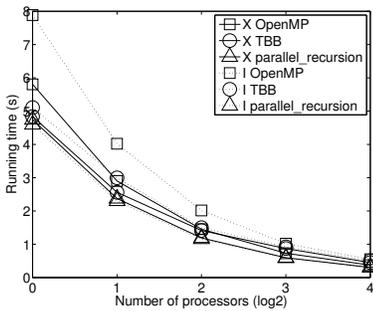


Figure 10. Performance of fib

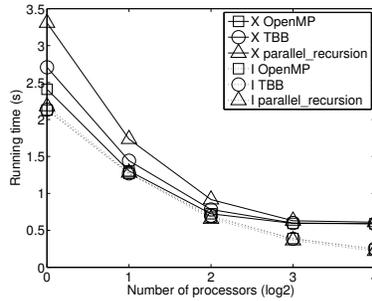


Figure 11. Performance of merge

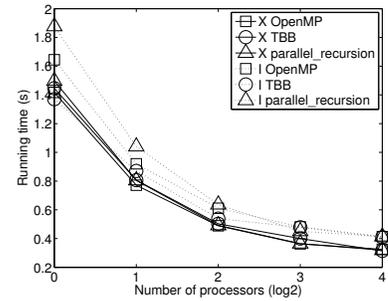


Figure 12. Performance of quicksort

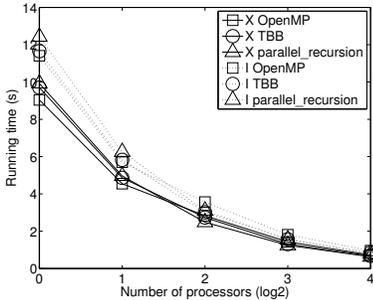


Figure 13. Performance of nqueens

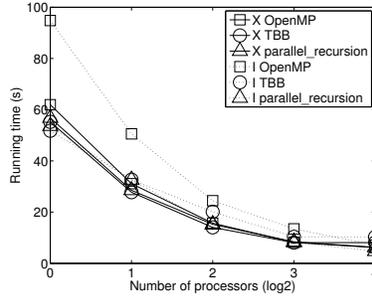


Figure 14. Performance of treadd

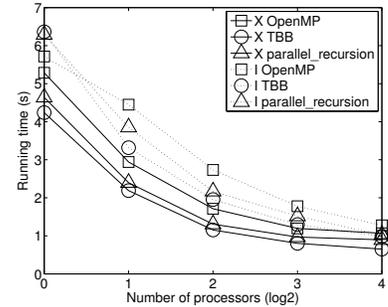


Figure 15. Performance of bisort

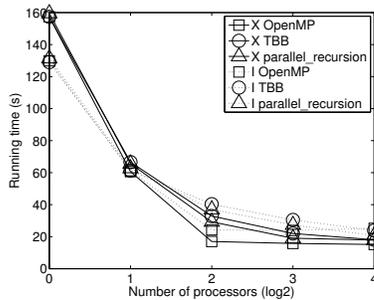


Figure 16. Performance of health

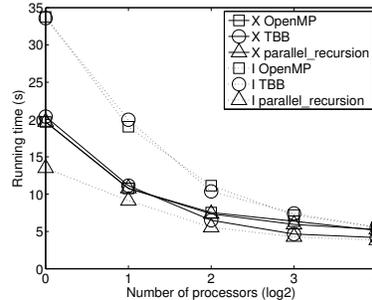


Figure 17. Performance of tsp

The OpenMP versions require conditionals to decide when to spawn new tasks or proceed sequentially. We found that including the conditional within the compiler directive led to unacceptable slowdowns. This way, the OpenMP codes with performance problems were written with a conditional in the code (not in a directive) that chooses between a code annotated with parallelization directives and a version without directives. Despite our attempts to tune these conditionals, their evaluation in each invocation of the recursive functions can be an important overhead for codes like fib and treadd whose computations are very simple. Another option would have been to write a version of the function without conditionals to be used when a level of recursion is reached in which the execution will not generate new parallel subtasks. This is similar to the existence of functions such as `Fib :: fib` in Figure 1 or the `TreeAdd` function invoked by the body in Figure 2, but while using these functions is the natural approach for the TBB templates, this is not the case for OpenMP. Also, let us remember that `parallel_recursion` removes the need for those functions. For these reasons we think this is the fairest comparison.

Since `parallel_recursion` is built on top of the TBB one could expect its codes to be slower than those based on `parallel_for` or `parallel_reduce`. This is not the case because our template is built directly on the low level task API provided by the TBB, not on its top level algorithm templates. Also, it follows different policies to

decide to spawn tasks and has different synchronization and data structure support requirements as we have seen. This makes it possible for `parallel_recursion` to be competitive with the native TBB version, and even win systematically in some benchmarks like `fib`. In other benchmarks like `merge` in the Xeon or `quicksort` in the Itanium `parallel_recursion` is non negligibly slower than the standard TBB when few cores are used, but the difference vanishes as the number of cores increases. The slowdowns of `parallel_recursion` in these two codes are due to operations repeated in each invocation to `child` to generate a subproblem. Generating at once a vector with all the children tasks could avoid this. Evaluating this option is part of our future work. The behavior of `tsp` in the Itanium is due to the compiler, as with `g++ 4.1.2` with the same flags the performance of all the implementations is very similar. Over all the benchmarks and numbers of cores, on average `parallel_recursion` is 0.3% and 19.7% faster than the TBB algorithm templates in the Xeon and the Itanium, respectively. If `tsp` is ignored in the Itanium due to its strange behavior, `parallel_recursion` advantage is still 9% in this platform. Its speedups over OpenMP are 2.5% and 30.5% in the Xeon and the Itanium, respectively; 21.4% in the latter without `tsp`.

Finally, we also experimented with trying to improve the performance of the template-based approaches using the partitioners that allow to control manually the subdivision of tasks in the runs with 16 cores. With the best parameters we found, the version based on standard TBB was on average 6.5% faster than the `parallel_recursion` based one in the Xeon, while `parallel_recursion` continued to lead the performance in the Itanium platform by 8%, or 2.4% if `tsp` is not counted.

## VI. RELATED WORK

While TBB is probably the most widespread library of skeletal operations nowadays, partially thanks to its support by Intel, it is not by any means the only one. The `eSkel` library [8] offers parallel skeletal operations for C on top of MPI. Its API is somewhat low-level, with many MPI-specific implementation details. Since C is not object oriented, it cannot exploit the advantages of objects for encapsulation, polymorphism, and generic programming where available, as is the case of C++. A step in this direction was Muesli [6], which is also oriented to distributed memory, being centered around distributed containers and skeleton classes that define process topologies. Muesli relies on runtime polymorphic calls, which generate potentially large overheads. This way [14] reports 20% to 100% overheads for simple applications. Lithium [2] is a Java library oriented to distributed memory that exploits a macro data flow implementation schema instead of the more usual implementation templates, but it also relies extensively on runtime polymorphism. Quaff [9] avoids this following the same approach as the TBB and our proposal, namely relying on C++ template metaprogramming to resolve polymorphism at compile time. Quaff's most distinctive feature is that it leads the programmer to encode the task graph of the application by means of type definitions which are processed at compile time to produce optimized message-passing code. As a result, while it allows skeleton nesting, this nesting must be statically defined, just as type definitions must be. Thus tasks cannot be generated dynamically at arbitrary levels of recursion and problem subdivision as the TBBs do. This is quite sensible, since Quaff works on top of MPI, being mostly oriented to distributed memory systems. For this reason its `scm` (split-compute-merge) skeleton, which is the most appropriate one to express divide-and-conquer algorithms differs substantially from the TBB standard algorithm templates and `parallel_recursion`.

Beyond libraries, the divide-and-conquer pattern of parallelism has also been object of interest in languages specialized on it such as HDC [12], a subset of the Haskell functional language.

## VII. CONCLUSIONS

Divide-and-conquer is a powerful pattern of parallelism that is particularly suitable to structure parallel applications in multicore systems. In fact recent popular tools such as TBB make extensive usage of it to express

parallelism because of its design around recursion. Still, the skeletal operations provided by this library are not well suited for this pattern of parallelism in many situations, leading to poor programmability. The reasons for this problematic have been explained along this paper, illustrating them with examples. An analysis of the weak points of the TBB algorithm templates for this pattern has led us to design a new algorithm template that overcomes them. We implemented this template on top for the task API of the TBB so that it is compatible with all the TBB library and other threading packages, including OpenMP, and it benefits from the load balancing of the TBB scheduler. Additionally, our implementation uses template metaprogramming very much as the standard TBB in order to provide efficient polymorphism resolved at compile time.

The examples in the paper and an evaluation using several productivity measures indicate that our algorithm template, called `parallel_recursion`, can improve substantially the programmer productivity when expressing divide-and-conquer parallel problems using skeletons. The only exception are algorithms based on vectors and matrices which require no reduction, for which the TBB `parallel_for` algorithm template may be more suitable. As for performance, `parallel_recursion` is on average somewhat faster than the TBB templates when automatic partitioning is used to balance load. There is not a clear winner when we try to adjust manually the granularity of the parallel tasks.

As future work, we want to evaluate small variations in the interface and to develop an extension that is suitable for distributed memory systems. This will possibly imply adding notations of locality and proximity to data in order to optimize the data movements required.

## VIII. ACKNOWLEDGEMENTS

This work was supported by the Xunta de Galicia under project INCITE08PXIB105161PR and the Ministry of Science and Innovation, cofunded by the FEDER funds of the European Union, under the grant TIN2007-67536-C03-02. We also want to acknowledge the Centro de Supercomputación de Galicia (CESGA) for the usage of its supercomputers.

## REFERENCES

- [1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman, *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1974.
- [2] M. Aldinucci, M. Danelutto, and P. Teti, “An advanced environment supporting structured parallel programming in Java,” *Future Gener. Comput. Syst.*, vol. 19, no. 5, pp. 611–626, 2003.
- [3] P. Bientinesi, J. A. Gunnels, M. E. Myers, E. S. Quintana-Ortí, and R. A. van de Geijn, “The Science of Deriving Dense Linear Algebra Algorithms,” *ACM Trans. Math. Softw.*, vol. 31, no. 1, pp. 1–26, Mar. 2005.
- [4] D. R. Butenhof, *Programming with POSIX Threads*. Addison Wesley, 1997.
- [5] R. Chandra, L. Dagum, D. Kohr, D. Maydan, J. McDonald, and R. Menon, *Parallel programming in OpenMP*. Morgan Kaufmann Publishers Inc., 2001.
- [6] P. Ciechanowicz, M. Poldner, and H. Kuchen, “The Münster Skeleton Library Muesli - A Comprehensive Overview,” Univ. of Münster, Tech. Rep. Working Papers, ERCIS No. 7, 2009.
- [7] M. Cole, *Algorithmic skeletons: structured management of parallel computation*. MIT Press, 1991.
- [8] —, “Bringing skeletons out of the closet: a pragmatic manifesto for skeletal parallel programming,” *Parallel Computing*, vol. 30, no. 3, pp. 389–406, 2004.
- [9] J. Falcou, J. Sérot, T. Chateau, and J.-T. Lapresté, “Quaff: efficient C++ design for parallel skeletons,” *Parallel Computing*, vol. 32, no. 7-8, pp. 604–615, 2006.
- [10] M. Frigo, C. E. Leiserson, H. Prokop, and S. Ramachandran, “Cache-oblivious algorithms,” in *FOCS '99: Procs. 40th Annual Symp. on Foundations of Computer Science*, 1999, p. 285.
- [11] M. H. Halstead, *Elements of Software Science*. Elsevier, 1977.
- [12] C. A. Herrmann and C. Lengauer, “HDC: A Higher-Order Language for Divide-and-Conquer,” *Parallel Processing Letters*, vol. 10, no. 2/3, pp. 239–250, 2000.

- [13] R. M. Karp, "Probabilistic analysis of partitioning algorithms for the traveling-salesman problem in the plane." *Math. of Operations Research*, vol. 2, no. 3, pp. 209–224, 1977.
- [14] H. Kuchen, "A skeleton library," in *Proc. 8th Intl. Euro-Par Conf. on Parallel Processing*, 2002, pp. 620–629.
- [15] T. Mattson, B. Sanders, and B. Massingill, *Patterns for parallel programming*. Addison-Wesley Professional, 2004.
- [16] McCabe, "A Complexity Measure," *IEEE Transactions on Software Engineering*, vol. 2, pp. 308–320, 1976.
- [17] OpenMP Architecture Review Board, "OpenMP Program Interface Version 3.0," May 2008.
- [18] J. Reinders, *Intel Threading Building Blocks: Outfitting C++ for Multi-core Processor Parallelism*. O'Reilly, July 2007.
- [19] A. Rogers, M. C. Carlisle, J. H. Reppy, and L. J. Hendren, "Supporting Dynamic Data Structures on Distributed-Memory Machines," *ACM Trans. on Programming Languages and Systems*, vol. 17, no. 2, pp. 233–263, March 1995.