AperTO - Archivio Istituzionale Open Access dell'Università di Torino

RPL: A Domain-Specific Language for Designing and Implementing Parallel C++ Applications

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1611663> since 2016-11-14T15:28:40Z

Publisher:

IEEE

Published version:

DOI:10.1109/PDP.2016.122

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

RPL: A Domain-Specific Language for Designing and Implementing Parallel C++ Applications

V. Janjic, C. Brown, K. MacKenzie
and K. Hammond
Uni. of St Andrews, UK.

{vj, cmb21, kwm2, kh}@st-andrews.ac.uk

M. Danelutto
Uni. of Pisa
Italy.

marco.danelutto@unipi.it

M. Aldinucci
Uni. of Torino
Italy.

aldinuc@di.unito.it

J. Daniel Garcia
Uni. Carlos III Madrid
Spain.

josedaniel.garcia@uc3m.es

Abstract—Parallelising sequential applications is usually a very hard job, due to many different ways in which an application can be parallelised and a large number of programming models (each with its own advantages and disadvantages) that can be used. In this paper, we describe a method to semi-automatically generate and evaluate different parallelisations of the same application, allowing programmers to find the best parallelisation without significant manual reengineering of the code. We describe a novel, high-level domain-specific language, *Refactoring Pattern Language (RPL)*, that is used to represent the parallel structure of an application and to capture its extra-functional properties (such as service time). We then describe a set of RPL rewrite rules that can be used to generate alternative, but semantically equivalent, parallel structures (parallelisations) of the same application. We also describe the *RPL Shell* that can be used to evaluate these parallelisations, in terms of the desired extra-functional properties. Finally, we describe a set of C++ refactorings, targeting OpenMP, Intel TBB and FastFlow parallel programming models, that semi-automatically apply the desired parallelisation to the application's source code, therefore giving a parallel version of the code. We demonstrate how the RPL and the refactoring rules can be used to derive efficient parallelisations of two realistic C++ use cases (Image Convolution and Ant Colony Optimisation).

I. INTRODUCTION

Despite the emergence of multi-core and many-core systems, parallel programming is still a very laborious task that is difficult to get right. Most current application designers and programmers are not experts in writing parallel programs. Knowing *where* and *when* to introduce parallel constructs, as well as *what* construct to introduce, can be a daunting and seemingly *ad-hoc* process. As a result, parallelism is often introduced using an abundance of low-level concurrency primitives, such as explicit threading mechanisms and communication, which typically do not scale well and can lead to deadlock, race conditions etc. Furthermore, software engineering tasks such as porting to other parallel platforms and general code maintenance can then require huge efforts and often rely on good parallel systems expertise.

Parallel patterns attempt to hide away this complexity by providing parameterised implementations of common types of parallel operations, such as parallel *farm* (a.k.a. *parallel-for*), *pipeline* and *reduce*. While making the task of parallel programming easier, the programmer still needs to choose the appropriate pattern *structure* for his problem and instantiate the patterns properly; this is a highly complex process. In addition, there is a wide choice of pattern libraries for C++, with over 10 libraries available at the time of writing. Moreover, each of

Fig. 1. A process of parallelising C++ applications using RPL and refactoring.

these pattern libraries offers different benefits and limitations, including different sets of supported skeletons and primitives, and different complexities of writing parallel code. Porting sequential C++ code to use one of these libraries and then porting the parallel code to an alternative pattern library still requires very significant effort.

In this paper (illustrated in Figure 1) we exploit *refactoring* tool-support to offer semi-automatic program transformations that, i) transform sequential code into its parallel equivalent, by introducing all of the complex parallel implementation via skeleton libraries; and, ii) transform a parallel implementation written in one skeleton library into an alternative implementation using a different library. We also describe a new high-level domain-specific language, *Refactoring Pattern Language (RPL)*, for describing the abstract and high-level parallel structure of a C++ application, together with a system of *rewrite rules* that operate over the RPL expressions, allowing transformations between equivalent parallel structures of the same application. Furthermore, the rewrite system also provides runtime estimations, guiding the user to choose the optimal parallel structure. Finally, once an optimal structure is obtained, the rewrite system returns a *refactoring recipe*: a report that shows the steps necessary for the user to

then refactor their application using tool-support in order to introduce the correct skeletons and nestings to obtain this optimal configuration. This allows programmers to write, with minimal effort, performance-portable parallel applications and to adapt them easily to new architectures by changing between different parallel structures and implementations. Finally, we evaluate our approach by parallelising two realistic C++ applications, Image Convolution and Ant Colony Optimisation, demonstrating that we can achieve scalable speedups with small programming effort. To show generality, our approach is demonstrated on three parallel frameworks – FastFlow [2], Intel Threading Building Blocks (TBB [23]) and OpenMP [7].

The paper makes the following research contributions:

- 1) We introduce new and novel refactorings that semi-automatically (i.e., under user guidance) introduce parallelism into the sequential application and also adapt the existing parallel code to alternative parallel models, including FastFlow, OpenMP and Intel TBB;
- 2) We introduce a new (external) domain specific language (DSL) for describing parallel structure of applications, together with a system of rewrite rules that allow automatic generation of alternative parallel structures that capture the same (sequential) code;
- 3) We introduce a scripting language for the refactorings generated by the rewritings from (2);
- 4) We show how applications can be parallelised “from scratch” using our DSL and refactoring system; and,
- 5) We evaluate our parallelisation method on two realistic, showing how the user can quickly and easily scale his/her application by switching to an alternative parallel model using our refactoring system.

II. PATTERNS AND SKELETONS

In our approach, parallelism is described in terms of high-level parallel patterns, which describe common types of parallel operations while abstracting over concrete implementation details. In this paper, we restrict ourselves to three classical parallel patterns, which we consider to be among the most useful and most common:

- The *map* pattern represent computations where the single function, f , is applied to a set of independent inputs, $x_1, \dots, x_m$, all of which are readily available. The parallelism arises from applying the same function to different inputs in parallel.
- The *farm* pattern is similar to map, with the difference that the inputs to a farm are not readily available, but are instead obtained from some stream in a sequence.
- The *pipeline* pattern models a parallel pipeline, where a sequence of functions, $f_1, f_2, \dots, f_m$ are applied, in turn, to a sequence of independent inputs, $x_1, x_2, \dots, x_n$, appearing one after the other. The output of f_i becomes the input to f_{i+1} , so that parallelism arises from executing $f_{i+1}(f_i(\dots f_1(x_k) \dots))$ in parallel with $f_i(f_{i-1}(\dots f_1(x_{k+1}) \dots))$.

The actual implementation of these parallel patterns is expressed in terms of *algorithmic skeletons* [8], parametric parallel code artifacts (e.g. functions or classes) that are instantiated with application code and other key information

to give concrete parallel programs. Skeletons abstract away low-level complexities such as thread creation, communication, synchronisation, and load balancing.

A. Common Skeleton Libraries

In this paper, we consider three state-of-the-art parallel skeleton libraries: FastFlow[2], Intel TBB[23] and OpenMP[7]. In addition to showing the generality of our approach, our choice of the skeleton libraries is motivated by subtle differences that exist between them, making some of them more suitable for certain classes of problems than others. *FastFlow* supports *stream-parallelism* and unbounded data streams as first-class concepts, making it suitable for the cases where the input is fed continuously to the application (e.g. video-stream processing). *OpenMP* supports fewer skeletons out-of-box compared to FastFlow and Intel TBB, making some of them (most notably pipeline) harder to implement properly, but also supports general low-level concepts (tasks and task synchronisation) that enable implementing a wider class of parallel applications and skeletons. Finally, *Intel TBB* supports a wider class of skeletons out of the box than does OpenMP, including true stream-parallel patterns such as pipeline and farm (called *parallel for* in TBB), but does not provide constructs to extend or configure the set of provided skeletons.

III. THE REFACTORING PATTERN LANGUAGE (RPL)

The *Refactoring Pattern Language* (RPL) allows application programmers to design patterned applications without having to deal with low-level implementation details. RPL provides a number of different features, including:

- i) *efficient and concise representation of the full set of patterns used to express the parallel structure of applications*, allowing the application programmer to abstract the parallelism exploitation strategies from the actual code and to implement these strategies in terms of the available library mechanisms;
- ii) *support for pattern rewriting techniques, together with mechanisms to estimate the non-functional properties exposed by patterns*, allowing the application programmer to experiment with different patterns/pattern compositions at design time, right before going to the actual coding phase;

RPL has been designed as a high-level DSL, independent of the underlying programming language or tools used to implement the parallel applications. It provides the programmers with the ability to represent **pattern expressions**. RPL currently supports a limited set of stream- and data-parallel patterns, including those discussed in Sec. II, and supports the definition of pattern expressions via the following (simplified) grammar:

$$\text{Pat} ::= \text{seq} \mid \text{streamGen} \mid \text{streamDrain} \mid \text{pipe}(\text{Pat}, \text{Pat}) \mid \text{comp}(\text{Pat}, \text{Pat}) \mid \text{farm}(\text{Pat}) \mid \text{map}(\text{Pat}) \quad (1)$$

The *seq* pattern represents a sequential code wrapped into a function that accepts inputs of type T_{IN} and produces the results of type T_{OUT} . *streamGen* and *streamDrain* represent sequential patterns producing and consuming a stream of some data type T_{STREAM} . *comp* represents a sequential composition of two patterns, where first all outputs for

Rule	Name
$\text{seq}(P) \rightarrow \text{farm}(P)$	<code>farm_intro</code>
$\text{farm}(P) \rightarrow P$	<code>farm_elim</code>
$\text{comp}(P_1, P_2) \rightarrow \text{pipe}(P_1, P_2)$	<code>pipe_intro</code>
$\text{pipe}(P_1, P_2) \rightarrow \text{comp}(P_1, P_2)$	<code>pipe_elim</code>
$\text{comp}(\text{map}(P_1), \text{map}(P_2)) \rightarrow \text{map}(\text{comp}(P_1, P_2))$	<code>map_prom</code>
$\text{map}(P) \rightarrow \text{comp}(\text{streamGen}, \text{farm}(P), \text{streamDrain})$	<code>data_strm</code>

Fig. 2. Pattern rewriting rules

the first pattern are produced and then sent as an input to the second pattern. `pipe`, `farm` and `map` represent pipeline, farm and map patterns, as described in Section II. As an example, a video processing application applying a set of filters to video frames may be expressed in RPL as `pipe(streamGen, seq(filt1), ..., seq(filtn), streamDrain)`.

Each term in a pattern expression can have one or more **attributes** (which represent different extra-functional properties of the pattern) associated with it. We will denote this by $\text{Pat} \leftarrow \text{attribute}$. For example, we say that the term has a *parallelism degree* of n (denoted by $\text{NW}(n)$) if n units of work (e.g. operating system threads) are used in parallel for the execution of that term. We can compute the parallelism degree of a pattern expression using the following simple rules:

- $\text{seq} \leftarrow \text{NW}(1)$;
- parallel degrees of farm and seq patterns are assigned by a programmer (or programming tool);
- if $\text{Pat1} \leftarrow \text{NW}(m)$ and $\text{Pat2} \leftarrow \text{NW}(n)$, then $\text{comp}(\text{Pat1}, \text{Pat2}) \leftarrow \text{NW}(\max(m, n))$;
- if $\text{Pat1} \leftarrow \text{NW}(m)$ and $\text{Pat2} \leftarrow \text{NW}(n)$, then $\text{pipe}(\text{Pat1}, \text{Pat2}) \leftarrow \text{NW}(m + n)$;

As another example, *service time* can be used in streaming applications as a performance measure. The service time of a stream pattern expression may be computed from the service times of its sequential components and of the parallelism degrees in the different patterns appearing in the expression [1]. Therefore, provided that we have obtained the service times of the seq nodes (e.g. using profiling), we can automatically calculate the service time of the whole pattern expression. We can use this information to compare different parallelisations and estimate which of them will give the best performance. We note that RPL focusses purely on specifying and reasoning about patterns, and, as such, we demonstrate RPL here using simple models for the sake of conciseness.

We have implemented a prototype of the *RPL shell* that provides the main features of the RPL described in this paper. The shell contains a set of well-known *pattern rewriting rules* (see Figure 2) that are used to generate alternative pattern expressions with the same functional semantics, but possibly different extra-functional properties. Below is a typical session the application programmer may run:

- 1) *Represent application parallel structure through RPL.* Suppose the application we want to represent is a simple pipeline with two stages, where the first stage processes an item in k time units and the second stage processes an item in $3k$ time units. The programmer may represent the parallel structure of his/her application as follows:

```
# let stage1 = Seq("f");;
```

```
# let stage2 = Seq("g");;
# let app = Pipe(stage1, stage2);;
# pp app;;
pipe
  seq f
  seq g
```

- 2) *Annotate pattern expressions with attributes.* The user can provide the service times of the seq terms.

```
# let stage1 = Aseq("f", [Ts(4.0)]);;
# let stage2 = Aseq("g", [Ts(12.0)]);;
# let app1 = Apipe(stage1, stage2, []);;
# ppa app1;;
pipe
  seq f with Ts(4.000000)
  seq g with Ts(12.000000)
```

- 3) *Evaluate extra-functional properties of the pattern expression.* The application programmer may then compute the attributes (in this case, the degree of parallelism and service time) for the top-level pattern expression, `app1`:

```
# servicetime app1;;
- : float = 15.
# pardegree app1;;
- : int = 2
```

- 4) *Evaluate alternative pattern expressions.* The application programmer, unsatisfied with the extra-functional features of his/her pattern expression, may apply some rewriting rules and ask for the properties of the new expression:

```
# let app2 = rr_pipe_elim(app1);;
# ppa app2;;
comp
  seq f with Ts(4.000000)
  seq g with Ts(12.000000)
# servicetime app2;;
- : float = 20.
# let app3 = rr_farm_intro app2 10;;
# ppa app3;;
farm with Pd(10) Te(1.000000)
comp
  seq f with Ts(4.000000)
  seq g with Ts(12.000000)
# servicetime app3;;
- : float = 1.6
```

It is worth pointing out that a completely different pattern expression could have been derived that leads to a comparable service time, but with different total parallelism degree (resources needed):

```
# let farm1 = rr_farm_intro stage1 3;;
# let farm2 = rr_farm_intro stage2 8;;
# let app4 = Apipe(farm1, farm2, []);;
# ppa app4;;
pipe
  farm with Pd(3) Te(1.000000)
  seq f with Ts(4.000000)
  farm with Pd(8) Te(1.000000)
  seq g with Ts(12.000000)
# servicetime app4;;
- : float = 1.5
# pardegree app3;;
- : int = 10
# pardegree app4;;
- : int = 11
```

- 5) *Generating a refactoring recipe* Once the programmer is satisfied with the extra-functional properties of the derived pattern expression, he/she can generate a *refactoring recipe* that contains the steps needed to transform the actual C++ application into its parallel version with the desired parallel structure. This recipe can then be used by the refactoring tool (Section IV) to parallelise the code. See Section V for example refactoring recipes.

We also have a prototype implementation of a *pattern expression rewriter*, that automates the process of generating

Fig. 3. Our refactoring tool built into Eclipse, complete with a menu of refactorings for FastFlow, OpenMP and TBB

alternative parallel structures and optimising extra-functional properties of interest. Based on the rewrite rules, it generates, from an initial pattern expression and using breadth-first search, all equivalent pattern expressions.

```
# let app_opt1 = ts_optim appl;;
# ppa app_opt1;;
pipe
  seq f with Ts(4.000000)
  farm with Pd(3)
    seq g with Ts(12.000000)
- : unit = ()
# servicetime app_opt1;;
- : float = 4.
```

The `ts_optim` rewriter in this case farms out the slower stages of the pipeline, so that the overall service time is reduced to the service time of the slowest stage.

IV. REFACTORINGS TO INTRODUCE PARALLEL PATTERNS

Refactoring is the process of changing the structure of a program, while preserving its functionality [19]. Unlike automatic program compilation and optimisation, refactoring emphasises the software development cycle, principally by: i) improving the design of software; ii) making software easier to understand; and iii) encouraging code reuse. This leads to increased productivity and programmability. Our refactoring tool (Figure 3) is implemented in the Eclipse development platform, using the CDT plugin for C++ development. This *semi-automatic* approach is more general than fully automated parallelisation techniques, which typically work only for a very limited set of cases under specific conditions, and are not easily tractable. Furthermore, unlike e.g. simple loop parallelisation, refactoring is applicable to a wide range of possible parallel structures, since parallelism is introduced in a structured way through algorithmic skeletons.

In this section we discuss new refactoring implementations which build upon earlier refactorings that introduce FastFlow skeletons [6]. We have implemented a set of refactorings for both Intel TBB and OpenMP. Due to space limitations, in this paper we focus on Intel TBB refactorings, but the OpenMP refactorings are similar.

We present two refactorings: *introduce map* and *introduce pipeline*. Each of these has two variants, according to the type of *component* involved. A *component* (or *task*) is a piece of code that will be executed in parallel as part of a parallel pattern. Typically components are computationally intensive pieces of code, such as expensive functions.

In TBB, components can be defined in two ways: using either *lambda expressions* or *function objects* (sometimes called *functors* in C++). These two approaches provide different advantages to the user. Lambda expressions (introduced in C++11 ([24, 11.4])) are anonymous functions which are useful for simple operations which are only used at a single point in the program and consist of only two or three statements. If the application requires the same operation in many places, or contains more complex code, functors are preferable: these are classes which overload the function application operator `()`, allowing instances of the class to be applied as if they were functions ([24, 3.4.3]).

```
1 for(int i = 0 ; i < NIMGS; i++) {
2   char *name = image_names[i];
3   image *img = read_image(name);
4   output[i] = filter (img);
5 }
```

Fig. 4. Sequential code before refactoring

We will illustrate our refactorings by showing the results when they are applied to the sequential code in Figure 4. This represents a loop which iterates over an array of strings containing names of image files: it reads each image from disk, applies some function called `filter`, and stores the result in the array `output`.

A. Parallel Maps

Maps using lambda expressions. This refactoring takes a loop, converts the body into a lambda expression, and applies TBB's `parallel_for` function to execute the iterations of the loop in parallel. If we apply this refactoring to the sequential code in Figure 4, the original loop is replaced with the code shown in Figure 5.

```
1 tbb::parallel_for(tbb::blocked_range<int>(0, NIMGS),
2   [&](const tbb::blocked_range<int>& r) {
3     for (int i = r.begin(); i != r.end(); ++i) {
4       char* name = image_names[i];
5       image* img = read_image(name);
6       output[i] = filter (img);
7     } });
```

Fig. 5. Code after introduction of a TBB lambda map

The result consists of a single expression which applies the `parallel_for` construct to two arguments: a TBB object called a *blocked range* which represents the range of values for which parallel execution is to take place, and a C++ lambda expression encapsulating the original sequential code. The syntax `[&]` at the start of the lambda expression indicates that any free variables in the body of the lambda expression should be captured by reference. It is also possible to use `[=]` to indicate capture-by-value, or `[]` to indicate that no variables are expected to be captured. When the refactoring is initiated, the user is presented with a dialogue where the capture method can be specified, and also the name of the parameter for the

blocked range object (`r` in Figure 5). The rest of the process is entirely automatic.

Maps using function objects. As an alternative to the lambda-based refactoring shown above, the user can choose to introduce a map based on a function object. In this case, a new class is introduced in the source code immediately before the function containing the code which is being refactored, and the original loop is replaced by a single `parallel_for` statement. The results for our sample code are shown in Figures 6 and 7.

```

1 class ProcessImage {
2     char** image_names;
3     char** output;
4
5 public:
6     void operator()(const tbb::blocked_range<int>& r) const {
7         for (int i = r.begin(); i != r.end(); ++i) {
8             char* name = image_names[i];
9             image* img = read_image(name);
10            output[i] = filter (img);
11        }
12    }
13
14    ProcessImage(char** arg_image_names, char** arg_output) :
15        image_names(arg_image_names), output(arg_output) {}
16 };

```

Fig. 6. C++ function object definition encapsulating sequential code

```

1 tbb::parallel_for(tbb::blocked_range<int>(0, NIMGS),
2                 ProcessImage(image_names, output));

```

Fig. 7. Invocation of TBB `parallel_for` statement to apply function object

As with the lambda map refactoring, the user is prompted for certain parameters (the name of the functor class and the name of the blocked range, `ProcessImage` and `r` in this case), but the rest of the process is automatic. The definition of the class `ProcessImage` is considerably more complex than the corresponding lambda expression in Figure 5 since the refactoring has to generate an entire class definition; moreover, since the class definition is in a different scope from the original code, the class definition has to include code to capture free variables from the original scope. In lines 2 and 3 of Figure 6 fields are defined which will contain pointers to the arrays `image_names` and `output` in the original code; these are filled in by the constructor on lines 14 and 15, which is applied in line 2 of Figure 7, at the site of the original code. Lines 5–12 in the class definition override the function call operator `()`, allowing `parallel_for` to apply instances of `ProcessImage` like functions. Note that writing class definitions like this by hand would be quite tedious and error-prone; our experience is that automatic refactoring makes it considerably easier to write parallel code.

B. Parallel Pipelines

We also have a refactoring to introduce pipelines using TBB constructs, again with two variants (lambda-based and functor-based). These require a `for` loop in which the body performs a chain of computations on the elements of an array, writing the eventual results to a possibly different array. Conveniently, the sequential code in Figure 4 is of this form, so

```

1 tbb::filter_t <void, char*> stage_1(tbb::filter::serial_in_order,
2                                   Closure_1(image_names));
3 tbb::filter_t <char*, image*> stage_2(tbb::filter::serial, Closure_2());
4 tbb::filter_t <image*, void> stage_3(tbb::filter::serial_in_order,
5                                   Closure_3(output));
6 tbb::parallel_pipeline(16, stage_1 & stage_2 & stage_3);

```

Fig. 8. Introducing a TBB pipeline

we can use it to demonstrate the refactorings. For the lambda-based refactoring, the refactored code is quite lengthy (25 lines), and we won’t discuss it in detail here.

Part of the result of the functor-based pipeline refactoring is shown in Figure 8. This is the code which replaces the original loop; in addition, three new classes are generated, one for each statement in the body of the loop. These are called `Closure_1`, `Closure_2` and `Closure_3`, and their definitions (omitted to save space) are quite similar to that in Figure 6.

The refactoring introduces three TBB objects (`stage_1`, `stage_2` and `stage_3`) of type `filter_t`, representing the three stages of the pipeline. In line 6 these are combined into a single pipeline using the `&` operator (overloaded by TBB) and executed using TBB’s `parallel_pipeline` function. The first parameter (16 here) represents the maximum number of stages of the pipeline which will be run in parallel. This value (which could be some more complex expression) is supplied by the user when the refactoring is initiated, as are the prefixes (`Closure` and `stage`) for the class definitions and stage names.

C. Remark on components

As can be seen from the above examples, the character of the refactored code is quite different depending on whether lambda expressions or function objects are used. Lambda expressions are quite concise and only require modification of the source code at the original site; however, the refactored code can be quite hard to read. In contrast, function objects lead to quite readable code at the expense of introducing large class definitions into the source code. There is however a possible advantage in that function objects have the potential to be re-used at multiple points. In our current TBB refactorings, component introduction happens automatically during other refactorings, but in future we may introduce a separate refactoring to introduce function objects. Our FastFlow refactorings already use this technique, so this would make the refactoring process more uniform; it would also give the user finer control over the structure of TBB refactorings.

V. EXPERIMENTS

In this section, we evaluate the effectiveness of the RPL (Section III) and the refactoring system (Section IV) for parallelising two realistic C++ use cases – Image Convolution and Ant Colony Optimisation. For each use case, we go through the following steps: i) starting from a sequential application, we abstract the structure into an RPL expression; ii) using the RPL shell and pattern rewriting rules from Figure 2, we generate alternative equivalent parallel structures (parallelisations) that have the same semantics as the original application and we pick those that have the best (estimated) runtime; iii) we

	<i>titanic</i>	<i>xookik</i>	<i>power</i>
Arch	AMD	Intel	IBM
Proc	Opteron 6176	Xeon x5675	Power8
Cores	24	6	20
Threads	24	12	160
Freq.	2.3 GHz	3.06 GHz	3.69 GHz
L2 Cache	24 x 512 Kb	12 x 256 Kb	20 x 512 Kb
L3 Cache	4 x 6 Mb	2 x 12 Mb	20 x 8 Mb
RAM	32 GB	48 GB	256GB

Fig. 9. Experimental Setup

generate *refactoring recipes* for parallel structures from the previous step; iv) for each recipe, we apply the corresponding refactorings (described in Section IV) to the sequential application, resulting in different parallel versions of the code; and, v) we evaluate the speedups (compared to the original sequential version) of the parallel versions from the previous step. To demonstrate the generality of our approach, for each recipe in step iv) we produce FastFlow, TBB and OpenMP versions of the original application (using refactorings from Section IV). Our evaluations of the speedups were conducted on three different parallel machines, *titanic*, *xookik* and *power*. Specifications of these machines are given in Figure V.

A. Image Convolution

Image convolution is a technique widely used in image processing applications for blurring, smoothing and edge detection. We consider an instance of the image convolution from video processing applications, where we are given a list of images that are first read from a file and then processed by applying a filter. Applying a filter to an image consists of computing a scalar product of the filter weights with the input pixels within a window surrounding each of the output pixels:

$$out(i, j) = \sum_m \sum_n in(i - n, j - m) \times filt(n, m) \quad (2)$$

The main hotspot for parallelisation is a **for** loop that iterates over the images, reading each one from a file and applying the filter:

```

1 for (int i=0; i<nr_images; i++) {
2   in_image[i] = read_image(i);
3   out_image[i] = process_image(in_image[i]);
4 }

```

Using the RPL syntax, we can abstract the structure of this code as

```
seq (comp (seq (read_image), seq (process_image))).
```

In order to parallelise the code, we can convert the **comp** into **pipe** and each of the three **seq** into **farm** (using, respectively, the *pipe_intro* and *farm_intro* transformations from Figure 2). There are, therefore, 15 (2^4 minus sequential version) different parallelisations of the above code. We will focus on the three that have the best estimated performance (via RPL shell):

```
1) seq (pipe (seq (read_image), map (process_image)))
```

```
2) seq (pipe (map (read_image), map (process_image)))
3) map (comp (seq (read_image), seq (process_image)))
```

Each of these parallelisations yields a different refactoring recipe. For example, parallelisation 2) leads to the following refactorings steps:

- 1) C_1 = Apply *identify-task* to *read_image*;
- 2) C_2 = Apply *identify-task* to *process_image*;
- 3) P_1 = Apply *introduce-pipeline* to *comp*;
- 4) F_1 = Apply *introduce-farm* to C_1 ;
- 5) F_2 = Apply *introduce-farm* to C_2 ;

Similar recipes can be derived for other parallelisations. We can then (semi-automatically) apply these recipes in order using our refactorings to get different parallelisations of the application. Figure 10 shows the speedups of different parallel versions of the application using FastFlow, TBB and OpenMP on *titanic* and *xookik*. In the figure, $(m \mid s)$ denotes the parallelisation where the **comp** is refactored into a **pipeline** and its first stage, **seq** (*read_image*), is refactored into a **map**; $(s \mid m)$ denotes the similar parallelisation, but where the second pipeline stage, **seq** (*process_image*), is rewritten into a **map** instead of the first; and, m denotes the parallelisation where the top-level **seq** is refactored into a **map**. We can observe similar results on all three architectures. When smaller number of threads is used, all parallelisations with all models perform similarly. For larger number of cores (e.g. 5 on *titanic*), $(s \mid m)$ parallelisation stops scaling with all models. m and $(m \mid m)$ parallelisations continue to scale well, reaching speedups of about 20, 10 and 32 on *titanic*, *xookik* and *power*, respectively. From these experiments, we can make three conclusions. Firstly, we are able to obtain very good speedups using our approach for this example. Secondly, the performance under all three libraries (OpenMP, TBB and FastFlow) is similar on all systems with all parallelisation. Therefore, how well the solution performs depends mostly on the way in which it is parallelised. Thirdly, it is crucial to choose the appropriate parallelisation, since the best parallelisations give significantly better speedups than others.

B. Ant Colony

Ant Colony Optimisation (ACO) [10] is a metaheuristic used for solving NP-hard combinatorial optimisation problems. In this paper, we apply ACO to the Single Machine Total Weighted Tardiness Problem (SMTWTP) optimisation problem, where we are given n jobs and each job, i , is characterised by its processing time, p_i (p in the code below), deadline, d_i (d in the code below), and weight, w_i (w in the code below). The goal is to find the schedule of jobs that minimises the total weighted *tardiness*, defined as $\sum w_i \cdot \max\{0, C_i - d_i\}$ where C_i is the completion time of the job, i . The ACO solution to the SMTWTP problem consists of a number of iterations, where in each iteration each ant independently computes a schedule, and is biased by a *pheromone trail* (τ in the code below). The pheromone trail is stronger along previously successful routes and is defined by a matrix, τ , where $\tau[i, j]$ is the preference of assigning job j to the i -th place in the schedule. After all ants compute their solution, the best solution is chosen as the “running best”; the pheromone trail is updated accordingly, and the next iteration is started. The main part of the program is given below:

Fig. 10. Image Convolution speedups on *titanic*, *xookik* and *power*. Here, $\mid$ is a parallel pipeline, m is a parallel map and s is a sequential stage.

```

1 for (j=0; j<num_iter; j++) {
2   for (i=0; i<num_ants; i++)
3     cost[i] = solve (i,p,d,w,t);
4   best_t = pick_best(&best_result);
5   for (i=0; i<n; i++)
6     t[i] = update(i, best_t, best_result);
7 }

```

Since `pick_best` in Line 4 cannot start until all of the ants have computed their solutions, and the `for` loop that updates `t` cannot start until `pick_best` finishes, we have implicit ordering in the code above. Therefore, the structure can be described in the RPL with:

seq (solve) ; pick_best ; seq (update)

where `;` denotes the ordering between computations. Due to an ordering between `solve`, `pick_best` and `update`, the only way to parallelise the sequential code is to convert `seq (solve)` and/or `seq (update)` into maps. Therefore, the possible parallelisations are:

- 1) map (solve) ; pick_best ; update
- 2) solve ; pick_best ; map (update)
- 3) map (solve) ; pick_best ; map (update)

Since `solve` dominates computing time, we are going to consider only parallelisations 1) and 3). Speedups for these two parallelisations, on *titanic*, *xookik* and *power*, with a varying number of CPU threads used, are given in Figure 11. In the figure, we denote map by m and a sequential stage by s . Therefore, $(m ; s ; s)$ denotes that `solve` is a parallel map, `pick_best` is sequential and `update` is also sequential. From the Figure 11, we can observe (similarly to the Image Convolution example) that speedups are similar for all parallel libraries. The only exception is Fastflow on *power*, which gives slightly better speedup than other libraries. Furthermore, both different parallelisations give approximately the same speedups, with the $(m ; s ; m)$ parallelisation using more resources (threads) altogether. This indicates that it is not always the best idea to parallelise everything that can be parallelised. Finally, we can note that none of the libraries is able to achieve linear speedups, and on each system speedups tail off after certain number of threads is used. This is due to a fact that a lot of data is shared between threads and data-access is slower for

cores that are farther from the data. The maximum speedups achieved are 12, 11 and 16 on *titanic*, *xookik* and *power*.

VI. RELATED WORK

Early work in refactoring has been described in [22]. A good survey (as of 2004) can be found in [17]. There has so far been only a limited amount of work on refactoring for parallelism [5]. In [4], a parallel refactoring methodology for Erlang programs, including a refactoring tool, is introduced for Skeletons in Erlang. Unlike the work presented here, the technique is limited to Erlang and does not evaluate reductions in development time. Other work on parallel refactoring has mostly considered loop parallelisation in Fortran [21] and Java [9]. However, these approaches are limited to concrete and simple structural changes (e.g. loop unrolling).

Parallel design patterns are provided as algorithmic skeletons in a number of different parallel programming frameworks [11] and several different authors advocated the massive usage of patterns for writing parallel applications [16], [20] after the well-known Berkeley report [3] indicated parallel design patterns as a viable way to solve the problems related to the development of parallel applications with traditional (low level) parallel programming frameworks.

In the algorithmic skeleton research frameworks, there is a lot of work on improving extra-functional features of parallel programs by using pattern rewriting rules [15], [1], [13]. We use these rules to support design space exploration in our system. Other authors use rewriting/refactoring to support efficient code generation from skeletons/patterns [12], which is a similar concept w.r.t. our approach. Finally, [14] proposed a “parallel” embedded DSL exploiting annotations, differently from what we use here, which is an external DSL. Other authors proposed to adopt DSL approaches to parallel programming [18], [25] similarly to what we propose here, although the DSL proposed is an embedded DSL and mostly aims at targeting heterogenous CPU/GPU hardware.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a high-level domain-specific language, Refactoring Pattern Language (RPL), that can be used to concisely and efficiently capture parallel patterns, and therefore describe the parallel structure of an application. RPL can also capture extra-functional parameters of patterns, such as

Fig. 11. Ant Colony Optimisation Speedups on *titanic*, *xookik* and *power*. Here, $;$ is normal function composition, m is parallel map and s is sequential stage.

the service time and parallelism degree. We also demonstrated how RPL can be used in the design stage of application development, to experiment with different parallel structures of the same application to obtain the best parallelisations. We described a set of refactorings that allow semi-automatic implementations of the desired parallel structure (described in RPL) to the sequential application code. To demonstrate generality, our refactorings target three different pattern/skeleton libraries – OpenMP, Intel TBB and FastFlow. Finally, we evaluated how the RPL and refactorings can be used to parallelise two realistic C++ use cases, Image Convolution and Ant Colony Optimisation, obtaining very good speedups on three different architectures – Intel, AMD and Power. As future work, we plan to extend RPL to support additional patterns, such as Reduce, Stencil and Divide-and-Conquer. We also plan to evaluate RPL and our refactorings on larger use-cases. We also plan to extend the prediction model in the RPL shell to use a more sophisticated method based on machine learning for estimating the values of the extra-functional attributes of the patterns.

Acknowledgements This work has been partially supported by EU FP7-ICT-2013-10 project REPORA (No. 609666) and EU H2020-ICT-2014-1 project RePhrase (No. 644235).

REFERENCES

- [1] M. Aldinucci and M. Danelutto. Stream Parallel Skeleton Optimization. In *Proc. PDCS 1999*, pages 955–962. IASTED/ACTA press, 1999.
- [2] M. Aldinucci, M. Danelutto, P. Kilpatrick, M. Meneghin, and M. Torquati. Accelerating Code on Multi-cores with FastFlow. In *Proc. Euro-Par 2011*, pages 170–181, 2011.
- [3] K. Asanovic, R. Bodik, J. Demmel, T. Keaveny, K. Keutzer, J. Kubiatowicz, N. Morgan, D. Patterson, K. Sen, J. Wawrzynek, D. Wessel, and K. Yelick. A view of the parallel computing landscape. *Commun. ACM*, 52(10):56–67, Oct. 2009.
- [4] C. Brown, M. Danelutto, K. Hammond, P. Kilpatrick, and A. Elliott. Cost-directed refactoring for parallel Erlang programs. *International Journal of Parallel Programming*, 42(4):564–582, 2014.
- [5] C. Brown, K. Hammond, M. Danelutto, P. Kilpatrick, H. Schnier, and T. Breddin. Paraphrasing: Generating Parallel Programs Using Refactoring. In *Proc. FMCO 2013*, volume 7542 of *Lecture Notes in Computer Science*, pages 237–256. Springer Berlin Heidelberg, 2013.
- [6] C. Brown, V. Janjic, K. Hammond, H. Schöner, K. Idrees, and C. W. Glass. Agricultural Reform: More Efficient Farming Using Advanced Parallel Refactoring Tools. In *Proc. PDP 2014*, pages 36–43, 2014.
- [7] B. Chapman, G. Jost, and R. v. d. Pas. *Using OpenMP: Portable Shared Memory Parallel Programming (Scientific and Engineering Computation)*. The MIT Press, 2007.
- [8] M. Cole. Bringing Skeletons Out of the Closet: A Pragmatic Manifesto for Skeletal Parallel Programming. *Par. Comp.*, 30(3):389–406, 2004.
- [9] D. Dig, J. Marrero, and M. D. Ernst. Refactoring Sequential Java Code for Concurrency via Concurrent Libraries. In *Proc. ICSE '09*, pages 397–407. IEEE Computer Society, 2009.
- [10] M. Dorigo and T. Stützle. *Ant Colony Optimization*. Bradford Company, Scituate, MA, USA, 2004.
- [11] H. González-Vélez and M. Leyton. A survey of algorithmic skeleton frameworks: High-level structured parallel programming enablers. *Softw. Pract. Exper.*, 40(12):1135–1160, Nov. 2010.
- [12] S. Gorlatch. Domain-specific optimizations of composed parallel components. In C. Lengauer, D. S. Batory, C. Consel, and M. Odersky, editors, *Domain-Specific Program Generation*, volume 3016 of *Lecture Notes in Computer Science*, pages 274–290. Springer, 2003.
- [13] S. Gorlatch, C. Wedler, and C. Lengauer. Optimization Rules for Programming with Collective Operations. In *Proc. IPPS/SPDP '99*, pages 492–499. IEEE Computer Society, 1999.
- [14] D. Griebler and L. Fernandes. Towards a domain-specific language for patterns-oriented parallel programming. In *Programming Languages*, volume 8129 of *Springer LNCS*, pages 105–119. Springer, 2013.
- [15] K. Matsuzaki, K. Kakehi, H. Iwasaki, Z. Hu, and Y. Akashi. A Fusion-Embedded Skeleton Library. In *Proc. Euro-Par 2004*, volume 3149 of *Springer LNCS*, pages 644–653. Springer Berlin Heidelberg, 2004.
- [16] T. Mattson, B. Sanders, and B. Massingill. *Patterns for Parallel Programming*. Addison-Wesley Professional, first edition, 2004.
- [17] T. Mens and T. Tourwé. A survey of software refactoring. *IEEE Trans. Softw. Eng.*, 30(2):126–139, Feb. 2004.
- [18] K. Olukotun. Beyond parallel programming with domain specific languages. In J. Moreira and J. R. Larus, editors, *Proc. PPOPP '14*, pages 179–180. ACM, 2014.
- [19] W. F. Opydyke. *Refactoring Object-oriented Frameworks*. PhD thesis, Champaign, IL, USA, 1992. UMI Order No. GAX93-05645.
- [20] J. L. Ortega-Arjona. *Patterns for Parallel Software Design*. Wiley Publishing, 1st edition, 2010.
- [21] J. Overbey, S. Xanthos, R. Johnson, and B. Foote. Refactorings for Fortran and high-performance computing. In *Proc. SE-HPCS '05*, pages 37–39. ACM, 2005.
- [22] H. Partsch and R. Steinbrüggen. Program transformation systems. *ACM Comput. Surv.*, 15(3):199–236, Sept. 1983.
- [23] J. Reinders. *Intel Threading Building Blocks*. O'Reilly & Associates, Inc., Sebastopol, CA, USA, first edition, 2007.
- [24] B. Stroustrup. *The C++ Programming Language*. Addison-Wesley Professional, 4th edition, 2013.
- [25] A. K. Sajeeth, K. J. Brown, H. Lee, T. Rompf, H. Chafi, M. Odersky, and K. Olukotun. Delite: A compiler architecture for performance-oriented embedded domain-specific languages. *ACM Trans. Embedded Comput. Syst.*, 13(4s):134:1–134:25, 2014.