



**QUEEN'S
UNIVERSITY
BELFAST**

The RePhrase Extended Pattern Set for Data Intensive Parallel Computing

Danelutto, M., Tiziano, D. M., Daniele, D. S., Mencagli, G., Torquati, M., Aldinucci, M., & Kilpatrick, P. (2017). The RePhrase Extended Pattern Set for Data Intensive Parallel Computing. *International Journal of Parallel Programming*. Advance online publication. <https://doi.org/10.1007/s10766-017-0540-z>

Published in:
International Journal of Parallel Programming

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights
© 2017 Springer. This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights
Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Open Access
This research has been made openly available by Queen's academics and its Open Research team. We would love to hear how access to this research benefits you. – Share your feedback with us: <http://go.qub.ac.uk/oa-feedback>

The RePhrase Extended Pattern Set for Data Intensive Parallel Computing

Marco Danelutto*, Tiziano De Matteis*,
Daniele De Sensi*, Gabriele Mencagli*,
Massimo Torquati*, Marco Aldinucci*,
Peter Kilpatrick^o

November 16, 2017

Abstract We discuss the extended parallel pattern set identified within the EU-funded project *RePhrase* as a candidate pattern set to support data intensive applications targeting heterogeneous architectures. The set has been designed to include three classes of pattern, namely i) core patterns, modelling common, not necessarily data intensive parallelism exploitation patterns, usually to be used in composition; ii) high level patterns, modelling common, complex and complete parallelism exploitation patterns; and iii) building block patterns, modelling the single components of data intensive applications, suitable for use-in composition-to implement patterns not covered by the core and high level patterns. We discuss the expressive power of the *RePhrase* extended pattern set and results illustrating the performances that may be achieved with the *FastFlow* implementation of the high level patterns.

Keywords Parallel design patterns, data intensive computing, stream computing, algorithmic skeletons

1 Introduction

We live in a world driven by information: electronic devices, manufacturing equipment and information systems produce data, either automatically or as a result of user interaction, and so applications for managing data intensive computations are becoming ever more important. At the same time, the notable improvement in the hardware available for data processing has prompted the development of new, highly demanding algorithms and applications.

In this scenario, *Data Intensive Computing* is gaining importance as a means of collecting, analysing and unveiling the knowledge that this data

University of Pisa*, Queen's University Belfast^o, University of Torino*
E-mail: {marcod, dematteis, desensi, mencagli, torquati}@di.unipi.it*,
p.kilpatrick@qub.ac.uk^o, aldinuc@di.unito.to*

encapsulates. Clearly, this possibility constitutes a valuable opportunity for many businesses and scientific applications.

However, the design, development and tuning of efficient data intensive applications still represents a very challenging task. Of necessity, these applications must be designed and implemented as parallel applications. In addition to the usual problems related to parallel computing, these applications also confront the programmer with the problem of efficiently managing considerable amounts of data, often available as streams dictating precise performance constraints.

Parallel design patterns have been identified as a viable mechanism to support parallel programmers in the difficult task of designing and implementing efficient and portable parallel applications [7,23]. Several existing and widely used programming frameworks provide the programmer of parallel applications with ready to use parallel patterns. Google mapreduce [13], Hadoop [31] and OpenMP basically provide a single pattern while Intel Thread Building Blocks [26] and Microsoft TPL [24] both provide a larger set of patterns. The programming frameworks developed as algorithmic skeleton programming frameworks include comprehensive sets of parallel patterns provided as ready to use programming abstractions: FastFlow [3,10], Muesli [19], SkePU [18], SkeTo [17]. In some cases the programming frameworks may be used to target different kinds of architectures. For example, Muesli targets workstation clusters, shared memory multicores and GPUs and FastFlow targets multicores, GPUs and provides partial support to target clusters of workstations as well as FPGAs (through TPC [20]). Although some of the pattern frameworks mentioned above have been explicitly designed to support data intensive applications, there is no broadly accepted definition of the set of patterns needed to support data intensive applications.

Within *RePhrase*, an EU H2020 funded project begun in April 2015, we aim to define a set of parallel patterns supporting the development of efficient data intensive applications on heterogeneous hardware platforms. In particular, we provide a set of parallel design patterns as ready to use programming abstractions fully compliant with standard C++ (11 and subsequent standard releases) paired with a set of tools to support pattern introduction in existing or new C++ code via refactoring and to check and ensure certain properties of the resulting parallel code. Fig. 1 summarises the overall approach of the *RePhrase* project.

In this paper we introduce the parallel pattern set identified within the *RePhrase* project. Our contribution consists in the presentation of a comprehensive parallel pattern set along with some preliminary results demonstrating the expressive power of the patterns, together with performance results achieved with a *FastFlow* implementation of these patterns. The rest of the paper is organised as follows: Sec. 2 introduces the data intensive computing paradigm. Sec. 3 describes the full set of patterns included in the *RePhrase* extended pattern set. Sec. 4 discusses the expressive power and usability of the pattern set. Finally, Sec. 5 presents some preliminary results from the pattern set implementation using *FastFlow*.

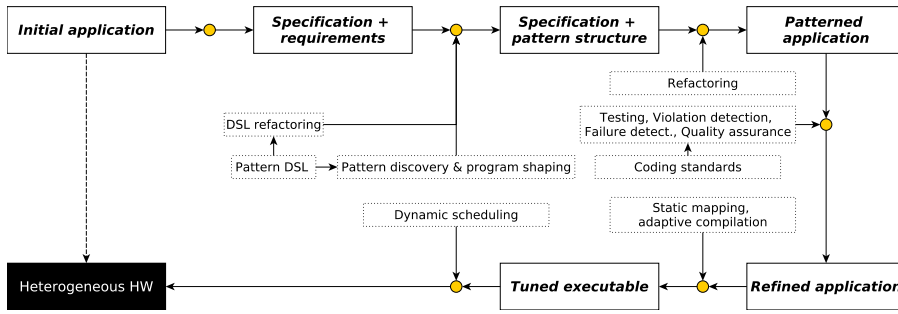


Fig. 1: *RePhrase* methodology workflow overview

2 Data Intensive Processing

As discussed in Sec. 1, *Big Data* is one of the top trending IT topics of today. It is characterised by the so-called 3Vs [21]: *variety*, *volume* and *velocity*. Variety refers to the nature and structure of the information. Volume refers to the magnitude of data produced. Finally, velocity refers to the frequency of data generation and to the dynamic aspects of the data in general. Different processing paradigms tackle various combinations of these aspects.

Pure *Data Parallel* systems tackle volume and variety: they process large masses of data, usually in an off-line fashion. Typically applications range across scientific sectors: examples include the analysis of massive data coming from scientific experiments [33], and studies of human digital traces (e.g. GPS traces) to discover and understand patterns in human mobility [25] or to support health care assistance [29]. Frameworks in this field take inspiration from Google’s *Map Reduce* [13]. Notable open source implementations include Apache Hadoop¹ and more recently Apache Spark², which is gaining attention due to its versatility and efficiency.

In turn, *Data Stream Processing* (DaSP) deals with the velocity and variety aspects of the “Big Data Challenge”. According to the DaSP paradigm, applications receive a *continuous* flow of data that has to be processed on the fly, usually with performance requirements in terms of bandwidth and/or latency [8,6]. Examples in this field include financial applications that try to spot revenue opportunities by analysing live market data [5], *Intrusion Detection Systems* that monitor network traffic in real-time to identify possible attacks [34], social media analytics that gather users’ news feeds and try to detect notable events [30]. Generally, applications are expressed as compositions of core functionalities in directed flow graphs, where vertices are *operators* (that encapsulate user defined logic) and arcs model streams, i.e. unbounded sequences of data items (*tuples*) sharing the same properties in terms of name

¹ <http://hadoop.apache.org/>

² <http://spark.apache.org/>

and type of attributes. Examples of solutions in this sphere include Apache Storm³, Apache Flink⁴ and IBM InfoSphere Stream⁵.

At times both aspects of data intensive processing can be present, allowing systems to serve a wider range of workloads and use cases. This approach is sometimes referred as *Lambda Architecture* [22]. In the *RePhrase* pattern set we include patterns that address both these aspects.

3 The *RePhrase* pattern set

Initially, the *RePhrase* set included two kinds of pattern:

- a set of *core* patterns, that comprises classical primitive parallelism exploitation patterns and may be specialised by means of a set of parameters to implement various applications using the pattern in slightly different ways; and
- a set of *high level* patterns, representing common, complex and specialised parallel patterns.

The first class includes, for example, pipeline and parallel for/map patterns, while the second includes examples such as divide&conquer and Google mapreduce patterns. Subsequently, taking into account the industrial use cases employed to assess the project results, we extended the pattern set with

- some further high level patterns, and
- with a collection of small *building block* patterns suitable for use, in composition, to model those data intensive patterns not captured by the *RePhrase* high level patterns.

It is worth pointing out that the “building block” patterns may, with relatively little effort, be used to implement the high level patterns in the *RePhrase* pattern set. However, due to the general purpose usefulness of these high level patterns, we preferred to provide them as primitives. This simplifies the implementation of parallel applications using the high level patterns and allows inclusion in their implementation of well-known optimisations, that would have been more difficult to include via the building block approach.

In the remainder of this section we introduce the patterns included in the *RePhrase* pattern set. The patterns are divided into classes according to the kind of parallelism exploited (data, task, stream, etc.). The description of the types of the patterns includes only the functional and code parameters, for the sake of simplicity.

³ <http://storm.apache.org/>

⁴ <http://flink.apache.org/>

⁵ <http://www-03.ibm.com/software/products/en/ibm-streams>

3.1 Stream Parallel “core” patterns

Stream parallel patterns employ parallelism in the processing of different items belonging to one or more input data streams. An input data stream is characterised by having a type⁶ and by being able to provide items (to be computed) one after the other with a given *inter-arrival time*. We will denote the type of a stream of data items of type α by α **stream**. A stream may be *finite*—in this case the last item of the stream will be the special item *eos*—or *infinite*. The infinite streams usually originate from an input device, e.g. a network card. Our core stream parallel patterns all process a single input stream to produce a single output stream.

Pipeline (**pipe**): the pattern computes in parallel several stages f_1, \dots, f_n on a stream of items, where $f_i : \alpha_{i-1} \rightarrow \alpha_i$ and $(\text{pipe } f_1 \dots f_n) : \alpha_0 \text{ stream} \rightarrow \alpha_n \text{ stream}$. Each stage processes data produced by the previous stage in the pipe and delivers results to the next stage. For each stream item x an item $f_n(f_{n-1}(\dots f_1(x)\dots))$ is eventually delivered in the pipeline output stream. Pipeline stages are executed in parallel.

Task-Farm (**farm**): the pattern computes in parallel a given function $f : \alpha \rightarrow \beta$ over all the items appearing in an input stream and so $(\text{farm } f) : \alpha \text{ stream} \rightarrow \beta \text{ stream}$ acts on input of type α **stream** delivering the results on the output stream of type β **stream**. Computations on different stream items are independent.

Stream Filter (**filter**): the pattern computes in parallel a filter $p : \alpha \rightarrow \{\text{true}, \text{false}\}$ over an input stream of type α **stream**, that is, it passes to the output stream only those input data items x such that $p(x) = \text{true}$. p must be a pure function and $(\text{filter } p) : \alpha \text{ stream} \rightarrow \alpha \text{ stream}$.

Stream Accumulator (**accumulator**): The pattern “sums up” using a binary function $\oplus : \alpha \times \alpha \rightarrow \alpha$ all items from the input stream and delivers the result to the output. The function used to sum up values (\oplus) may be any binary function of type $\oplus : \alpha \times \alpha \rightarrow \alpha$, although commutative and associative functions will provide much better and more scalable implementations. $(\text{accumulator } \oplus) : \alpha \text{ stream} \rightarrow \alpha$.

Stream Iteration (**iteration**): the pattern iterates the computation of another pattern over one or more items appearing on the input stream, and delivers results on the output stream. The pattern has type $\text{iteration } (\alpha \text{ stream} \rightarrow \alpha \text{ stream}) \times (\alpha \rightarrow \text{bool}) \rightarrow (\alpha \text{ stream} \rightarrow \alpha \text{ stream})$. The first parameter is the nested pattern, the second is the function used to redirect output item x to the input of the nested pattern (true) or to the output of the iteration pattern (false).

⁶ in the following we use Greek letters to denote data types. The expression $x : \alpha$ is used to denote an object x whose type is α while the expression $f : \alpha \rightarrow \beta$ is used to denote a function f computing a result of type β out of an input of type α .

3.2 Data Parallel “core” patterns

Data parallel patterns employ parallelism in the processing of different items or (possibly overlapping) partitions of items belonging to a single “collection” data item. The key point in this case is the existence of two (logical) functions decomposing a single input data collection (of type α collection) into a collection of collections (**decomp**: α collection \rightarrow (α collection) collection) and building the result out of the collection of subresults (**comp**: β collection \rightarrow γ). Data parallel patterns process a single collection at a time, but nothing prevents them being used to operate on a *stream* of collections to produce a *stream* of collections.

Map (**map**): this pattern computes a given function $f : \alpha \rightarrow \beta$ over all the data items of an input collection whose elements have type α (**map** $f : \alpha$ collection \rightarrow β collection). Therefore the **decomp** function (logically) returns a set of α singletons out of the α collection input and the **comp** rebuilds a β collection out of the collection of singleton results. Given the input collection x_1, \dots, x_N , the output collection is $f(x_1), \dots, f(x_N)$. Since each data item in the input collection is independent of the other items, all the elements can be computed in parallel.

Reduce (**reduce**): the pattern “sums up” all the data items of a collection of items of type α using a binary function $\oplus : \alpha \times \alpha \rightarrow \alpha$ which is usually associative and commutative (**reduce** $\oplus : \alpha$ collection \rightarrow α). Given the input collection x_1, \dots, x_N , the **reduce** computes $x_1 \oplus \dots \oplus x_n$.

Stencil (**stencil**): the pattern decomposes an input collection ($x : \alpha$ collection) to a set of as many sub collections ($y : \alpha$ collection) as the original collection component count. Each sub collection hosts a distinct item of the original collection along with a set of *neighbour* items. A function $f : \alpha$ collection \rightarrow β is used to compute in parallel the new values of the output $z : \beta$ collection.

3.3 High Level Patterns

High level patterns model more complex parallel computations. All are used to compute the result relative to a single input, although they may be used in composition with stream parallel patterns to compute a stream of results from a stream of inputs. We informally specify the intended parallel semantics.

Divide and Conquer (**dac**): the pattern computes a problem for which *a*) the solution for some base cases are known and *b*) non-base case problems may be divided into a collection of sub-problems and *c*) the solution of the non-base case problems may be computed from the solutions of the sub-problems. The type of the pattern is **dac** : $divide \times conquer \times isBaseCase \times SolveBaseCase \times \alpha \rightarrow \beta$ with $divide : (\alpha \rightarrow \alpha$ collection), $conquer : (\beta$ collection $\rightarrow \beta)$, $isBaseCase : (\alpha \rightarrow bool)$ and $solveBaseCase : (\alpha \rightarrow \beta)$

Mapreduce (**mapreduce**): the pattern computes the Google mapreduce [13] using two functions $f : \alpha \rightarrow \beta \times \kappa$ and $\oplus : \beta \times \beta \rightarrow \beta$, where κ is the key type, and has type **mapreduce** $f \oplus : \alpha \text{ collection} \rightarrow (\beta \times \kappa) \text{ collection}$. The first function (f) is used to map all the items in the input collection to $\langle \text{key}, \text{value} \rangle$ pairs, while the second (\oplus) is used to compute a unique value out of the *value* entries in $\langle \text{key}, \text{value} \rangle$ pairs with the same *key* value.

Pool pattern (**pool**): the pattern models the evolution of a population of individuals. Iteratively, selected individuals are subject to evolution steps. The resulting new individuals are inserted in the population or discarded according to their fitness score. The process is iterated up to a given number of iterations (or up to a given computation time) or up to the point that an individual with a given fitness is inserted in the population. Low fitness individuals may be removed from the population to keep the population size constant at each iteration. The type of this pattern is thus **pool** *sel evol fit merge term* : $\alpha \text{ collection} \rightarrow \alpha \text{ collection}$ where *sel* : $\alpha \text{ collection} \rightarrow \alpha \text{ collection}$, *evol* : $\alpha \rightarrow \alpha$, *fit* : $\alpha \rightarrow \beta$, *merge* : $\alpha \text{ collection} \times \alpha \text{ collection} \rightarrow \alpha \text{ collection}$, *term* : $\alpha \text{ collection} \rightarrow \text{bool}$.

Image convolution pattern (**convolve**): this pattern computes image convolution according to some input kernel parameter and has type **convolve** : $\alpha \text{ mat} \times \text{int mat} \rightarrow \alpha \text{ mat}$. A kernel parameter is an $N \times N$ matrix (usually 3×3 or 5×5) of integer values. The image convolution is obtained from the source image processing each pixel at position i, j by taking the $N \times N$ values centred at i, j , multiplying each of the values by the corresponding value of the kernel and summing up the results to get the new i, j pixel of the resulting matrix. Image convolution may be used to obtain different effects with different kernels, ranging from image blurring to image enhancement, embossing, sharpening, etc. The image convolution pattern may obviously be implemented using a stencil pattern, but it is provided as a first class pattern due to its wide usage.

Windowed stream farm (**windowedSF**): the pattern computes functions on windows of stream item values, and has type **windowedSF** : $\alpha \text{ stream} \times (\alpha \text{ vec} \rightarrow \beta) \rightarrow \beta \text{ stream}$. In particular, this pattern implements a computation that outputs items on the output stream corresponding to the evaluation of a given function over successive, consecutive windows of items appearing on the input stream. The windows have a length (number of items to be listed in the window) and an overlap factor (number of items in window w_i also appearing in window w_{i+1}). The number of items in a window may be defined either as an actual number (count-based windows) or as a time interval, that is, as the items appearing in the input stream within the given interval of time (time-based windows).

Keyed stream farm (**keyedSF**): the pattern computes functions on windows of stream item values, and has type **keyedSF** : α **stream** \times $(\alpha$ **vec** $\rightarrow \beta)$ \times $(\alpha \rightarrow \gamma$ **key**) $\rightarrow \beta$ **stream**. Each input item belongs to a unique class called key (with type γ **key**); that is, the physical stream can be viewed as a multiplexing of several logical streams, each of which conveys items with the same key value. This pattern implements a computation that outputs items on the output stream corresponding to the evaluation of a given function over successive, consecutive windows of items appearing on the same logical input stream. The windows have a length (number of items to be listed in the window) and an overlap factor (number of items in window w_i also appearing in window w_{i+1}). The number of items in a window may be defined either as an actual number (count-based windows) or as a time interval, that is as the items appearing onto the input stream within the given interval of time (time-based windows).

3.4 Data intensive building block patterns

The patterns in this class are further divided into patterns used to generate/collapse data streams and in patterns used to process existing streams.

3.4.1 Stream generate/collapse patterns

Stream generator pattern (**streamgen**): this pattern is used to generate a stream from an internal (e.g. a stateful function) or external (e.g. a disk file) data source and has type **streamgen** : $() \rightarrow \alpha$ **stream**⁷.

Stream collapser pattern (**streamdrain**): this pattern is used to “consume” all the items appearing on its input stream and has type **streamdrain** : α **stream** $\rightarrow ()$.

Data splitter pattern (**datasplitter**): the pattern is used to generate a stream of items out of the components of a data collection (possibly from the pattern input stream) according to a user-defined strategy and has type **datasplitter** : α **collection** $\rightarrow \beta$ **stream**, where β is either α or α **collection**.

Data merger pattern (**datamerger**): the pattern is used to gather items appearing onto an input stream in a data collection according to a user defined strategy and to deliver the data collection onto the pattern input stream and has type **datamerger** : α **stream** $\rightarrow (\alpha$ **collection**) **stream**.

3.4.2 Stream processing patterns

Stream filter pattern This is the same **filter** pattern as is included in the core stream patterns (Sec. 3.1. It is listed here as logically it belongs to the stream processing subclass of the data intensive building block patterns.

⁷ where $()$ is the “no parameter” (void) type

Stream merger pattern (**streamMerger**): this pattern is used to merge two or more input streams into a single output stream according to a pre-defined or user-specified merge policy and has type **streamMerger** : $(\alpha \text{ stream}) \text{ collection} \rightarrow \alpha \text{ stream}$.

Stream tupler pattern (**streamTupler**): this pattern processes items from a set of input streams to produce a tuple on a single output stream with exactly one item from each of the input streams and has type **streamTupler** : $\alpha_1 \text{ stream} \times \dots \times \alpha_m \text{ stream} \rightarrow (\alpha_1 \times \dots \times \alpha_m) \text{ stream}$.

Stream splitter pattern (**streamSplitter**): this pattern directs the items appearing on a single input stream to one of the different output streams according to a pre-defined or user-defined split policy and has type **streamMerger** : $\alpha \text{ stream} \rightarrow (\alpha \text{ stream}) \text{ collection}$.

Stream detupler pattern (**streamDetupler**): the pattern processes tuples appearing on an input stream. Each tuple is used to generate items on different output streams according to a parameter policy and has type **streamDetupler** : $(\alpha_1 \times \dots \times \alpha_m) \text{ stream} \rightarrow \alpha_1 \text{ stream} \times \dots \times \alpha_m \text{ stream}$. Default policies are provided including:

- scatter (tuple components to different output streams, in order)
- unicast (tuple components to the same output stream, one after the other, where the stream is identified through a user supplied function)

4 Expressive power of the *RePhrase* pattern set

We discuss the expressive power of the *RePhrase* extended pattern set from two viewpoints:

- the class of data intensive applications supported; and
- the programming effort required to code a data intensive application using the patterns in comparison with that required to program the applications using traditional, “non-patterned” programming frameworks.

4.1 Applications supported

The three different kinds of pattern provided within the *RePhrase* extended pattern set all support different, partially overlapping, classes of applications (see below).

We have recently implemented the Parsec⁸ benchmarks using **FastFlow** [12], the library we use to provide the application programmer with *RePhrase* patterns (see Fig. 4). We have also already confirmed the possibility of implementing all of the Cowichan problems [32] using the *RePhrase* pattern set. In addition, and obviously, the pattern set covers all the parallel needs of the *RePhrase* use case set [27].

⁸ <http://parsec.cs.princeton.edu/>

High level patterns Each of the high level patterns in the *RePhrase* extended pattern set supports a complex and complete set of well-know parallel patterns. In general each pattern may also be implemented using a (composition of) core pattern(s) although this alternative is not necessarily more efficient or easier to implement. As an example, a divide and conquer pattern may be implemented using a task farm pattern where the workers are able to compute all of the specific phases (divide, test base case, solve base case, conquer) and the tasks produced while dividing are routed back from collector to emitter for further processing. The implementation of the divide and conquer pattern in *FastFlow* follows a similar strategy, but implements a number of optimisations such that a high degree of efficiency is achieved in the implementation of a wide range of divide and conquer kernels and applications. The divide and conquer pattern thus supports the implementation of a variety of parallel algorithms, ranging from non-data intensive algorithms to data intensive ones including sorting of large datasets or computation bound algorithms such as Strassen dense matrix multiplication. The pool pattern is particularly suited to evolutionary computing applications. It has been demonstrated to be useful in the exploration of complex space search algorithms, in the implementation of genetic algorithm based applications and in the implementation of iterative algorithms modelling approximation of complex solutions through progressive refinement. Despite the fact we assume slightly different I/O mechanisms in the *RePhrase* pattern set, the Google mapreduce pattern has already been demonstrated useful in a number of different applications, and we do not spend more space here to justify its inclusion in the pattern set. Finally, the key and windowed stream farm patterns have been shown to efficiently support financial applications processing a data intensive stream of records and may naturally support those applications, e.g. from social networks, processing large sets of records available across single or multiple data streams to infer more structured information about the stream contents and behaviour.

Core patterns Core patterns, alone or in composition, may be used to support those applications and kernels where embarrassingly parallel, staged (i.e. pipelined) or iterative parallel components are present. They have been included in the *RePhrase* pattern set but have already been present in several other programming frameworks (including [17, 18, 19, 26, 10]). The class of applications supported includes numerical applications, video processing applications, soft computing and AI applications up to learning and massive data processing applications.

Building block patterns The building block patterns included in the *RePhrase* pattern set must be used in composition to model the parallel patterns needed for data intensive applications that are not supported by either high level or core patterns. As such, they provide support for the implementation of any generic streaming network built out of an arbitrary number of data sources and data drains with an arbitrary number of processing nodes, transforming,

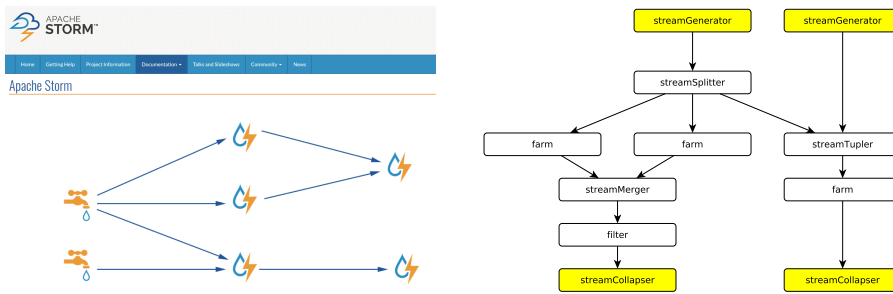


Fig. 2: Sample data streaming network with *RePhrase* building block and core patterns: Apache Storm website picture (left, from <http://storm.apache.org/>) and sample *RePhrase* building block outline (right).

filtering, merging, splitting and collapsing stream (portions). The set of patterns in the building block class have clearly been inspired by the kind of computations usually supported by programming frameworks such as Storm and Flink. A data processing network such as that in Fig. 2 may be easily built by combining our building block and core patterns.

4.2 Programming effort

The programming effort required to implement data intensive parallel applications varies according to the application at hand and to the targeted parallel programming framework.

A mapreduce application programmed on top of Hadoop simply requires specification of the code for the map and reduce “functions” along with some input data and the Hadoop framework turns these minimal inputs into an efficient, running application. In contrast, if you wish to program yourself the mapreduce pattern using MPI, the amount of code required increases substantially. On the other hand, if you wish to program a non-mapreduce application on top of Hadoop you may end up concluding that either this is not possible at all, or that the effort needed to mutate the mapreduce pattern into the actual pattern to be implemented is too great.

The programming effort required of the programmer using the *RePhrase* patterns is similar to that required of the Hadoop programmer developing a mapreduce application. *RePhrase* patterns are provided using plain C++ programming abstractions (higher order functions or classes) that may be instantiated with suitable functional and non functional parameters to implement the particular instance of the pattern required by the application programmer.

We discuss two simple examples here, relative to the use two of the main “technologies” adopted and developed within *RePhrase*: the *FastFlow* structured parallel programming environment and the *GrPPI* C++ pattern interface. *FastFlow* is one of the target back-ends considered within *RePhrase* and

```

// FASTFLOW two stage pipeline           // GrPPI two stage pipeline

auto f1 = [](T1 * x)->(T2*) { ... };      auto f1 = [](T1 x) { ... };
auto f2 = [](T2 * x) { ... };            auto f2 = [](T2 x) { ... };

struct Stage1 : ff_node_t<T1,T2> {
    T2 * svc(T1 * x) { return (f1(x)); }
};

struct Stage2 : ff_node_t<T2> {
    void * svc(T2 * x) { return(f2(x)); }
};

int main(int argc, char * argv[]) {      int main(int argc, char * argv[]) {
    ...                                    ...
    ff_Pipe pipe(Stage1,Stage2);           parallel_execution_ff ff_mode{};
    ...                                    pipeline(ff_mode, f1, f2);
    pipe.run_and_wait_end();              ...
    ...                                    }
}

```

Fig. 3: Pipeline sample code snippets in FastFlow and GrPPI

it is being currently developed and maintained at the Universities of Pisa and Torino⁹. GrPPI is a C++11 specific pattern interface designed within *RePhrase* to provide a target framework-agnostic way of expressing patterns [14]. It provides “functional style” patterns that can be used within standard C++ programs and uses several different back-ends to execute the pattern code in parallel, including OpenMP, Intel TBB and FastFlow. The overhead introduced by the GrPPI is negligible, w.r.t. to the usage of the back-end mechanisms to implement the same patterns, as GrPPI is provided as a header only library and its implementation adopts all the standard meta programming mechanisms, resulting in a large compile time effort but in a modest run time overhead.

Fig. 3 presents code snippets illustrating FastFlow and GrPPI pattern usage examples:

- In FastFlow a pipeline pattern with sequential stages may be expressed by declaring a `ff_pipeline` object and then adding stages (lambdas or `ff_node_t` objects wrapping a `function` object). Once the object has been declared, execution may be triggered by invoking the `run_and_wait_end()` method of the `ff_pipeline` object (see code snippet in Fig. 3 (left)).
- Using GrPPI a pipeline pattern may be declared and run as a function with parameters that denote the kind of target parallel programming framework and the stages to be used in the pipeline (this is a variable length list of `callable` object) (see code snippet in Fig. 3 (right)).

Overall, the extended pattern set provides significant support for the parallel applications programmer in the implementation process by making available the patterns as ready to use objects and functions that the programmer

⁹ <http://calvados.di.unipi.it/fastflow>




Fig. 4: FastFlow RePhrase patterns vs. standard programming frameworks: results on Parsec benchmarks. LOC (lines of code, top), code churn (number of changed lines, middle) and execution times on an Ivy Bridge Xeon server (bottom). (FF = FastFlow, TBB = Intel Thread Building, OMP = OpenMP, OMPSS = OpenMP SuperScalar [15])

may freely and immediately use to program the parallel part(s) of his/her application. However, it is worth noting that the *RePhrase* methodology (as depicted in Fig. 1) aims at introducing patterns into applications by means of the *RePhrase* refactoring tools. Places where patterns may be introduced are determined by using the *pattern discovery* tool which identifies those locations and those portions of code that may be turned into parallel pattern instances.

In terms of LOC (lines of code), the programming effort required to use the native FastFlow pattern interface is comparable to that required by similar programming frameworks (e.g. Intel TBB [26]) but certainly FastFlow requires a greater number of lines of code with respect to pragma based programming frameworks such as OpenMP. However, we must point out that the FastFlow

programming interface provides many more patterns than OpenMP. For those natively supported in OpenMP—e.g. `parallel for/map`—LOC is better (lower) in OpenMP, but those that are not natively supported in OpenMP require a comparable or even larger LOC in OpenMP compared to FastFlow (see Fig. 4). A completely different perspective comes from the use of GrPPI [14]. In this case the LOC count is balanced even when comparing the *RePhrase* framework with pragma based frameworks such as OpenMP, due to the fact the GrPPI profitably leverages those new features recently added to the C++ standard that *de facto* support functional style abstractions. It is worth pointing out that, while pragma based patterns require some intervention on the compiler toolchain, the *RePhrase* wrapper approach implemented in GrPPI works with (pre-compiled or source header only) libraries.

Targeting GPUs As far as GPU targeting is concerned, the *RePhrase* pattern set offers different possibilities. On the one hand, GrPPI will eventually provide a C++/Thrust implementation of those patterns that may be fruitfully executed on a GPU (e.g. data parallel patterns such as the `map` or `reduce` patterns) [14]. In this case, the only GPU specific change that will be needed in the code is the specification of the GPU specific `parallel_execution_thrust` parameter in the GrPPI pattern call (this is the substitute of the `ff_mode` parameter in the GrPPI code of Fig. 3 (right)). On the other hand, FastFlow offers the possibility to use GPUs through the main data parallel patterns, using minimal changes in the code. Consider a `map` pattern.

The typical changes required are just two (see Fig. 5). First, the function to be computed in the GPU kernel has to be specified using suitable macros to support on-the-fly kernel code generation. The macro has parameters such as the name of the function, the type of the input parameter, the name of that parameter and the code itself used to transform the parameter into the kernel result. Second, the programmer must supply an object with a `setTask` method specifying the pointers to the input and output data and the data size.

Use of the `map` requires code similar to that needed to call the same `map` targeting the CPU cores. Fig. 5 (right) shows a `main` using a `farm` pattern whose workers are `map` patterns executed on the GPU. The left part of the figure shows the GPU specific code needed to use the GPU `map`.

5 Experimental results

In this section we present some experimental results obtained with the high level patterns in the *RePhrase* extended pattern set. Most of the results are relative to the *native* FastFlow implementation of the patterns as GrPPI does not introduce a significant performance penalty.

```


FFMAPFUNC(mapF, unsigned int, in,      ...
           return in + 1;
);
class cudaTask: public
  baseCUDATask<unsigned int,
              unsigned int> {
public:
  void setTask(void* t) {
    if (t) {
      cudaTask *t_ = (cudaTask *)t;
      setIntPtr(t->in);
      setOutPtr(t->in);
      setOutPtr(t->out);
      setSizeIn(inputsize);
    }
  }
};
unsigned int *in, *out;
};
int main(int argc, char* argv[]) {
  ...
  ff_farm<> farm;
  Emitter E(streamlen,inputsize);
  Collector C(inputsize);
  farm.add_emitter(&E);
  farm.add_collector(&C);
  std::vector<ff_node *> w;
  for(int i=0;i<nworkers;++i)
    w.push_back(
      new FFMAPCUDA(cudaTask, mapF)());
  farm.add_workers(w);
  farm.run_and_wait_end();
  ...
};

```


Fig. 5: Targeting GPUs in FastFlow: Farm with GPU Map workers

5.1 Divide & Conquer

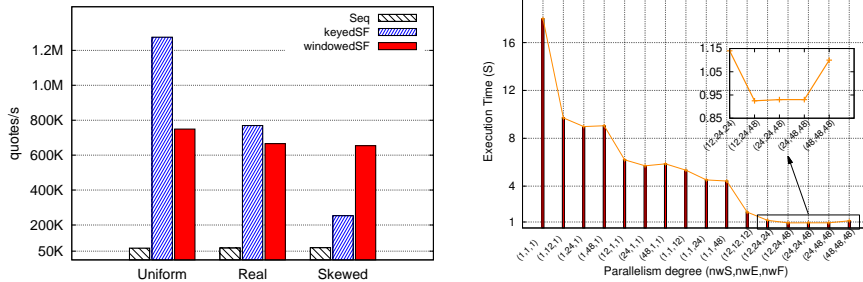
We implemented the Divide&Conquer (DAC) parallel pattern in various back-end environments, such that, while maintaining the same source code, the programmer can exploit the potential of different frameworks and target architectures (see Fig. 6a). We proposed three different implementations for multicore architectures based on OpenMP compiler annotations, Intel TBB and FastFlow parallel programming libraries. The experimental analysis, performed on a 24-core Intel server, showed that the reduced effort in programming does not come at the expense of significant performance penalties. The experimental study has been carried out by comparing the pattern-based solution with hand-made parallelisations using the same back-end runtime. These results pave the way for further development of this work. First, the set of back-end implementations can be further extended, including an MPI implementation for targeting distributed systems, and a CUDA/OpenCL-based implementation for GPUs. Second, we recognise that an important role in achieving a good level of performance is played by the cutoff value, i.e. the point at which we stop the recursion and solve the problem sequentially to better exploit the cache hierarchy and/or limit the runtime support overhead. This value depends on the structure of the specific parallelised application and on the kind of platform used. As proposed in [16], using information from the application collected at runtime (without relying on user hints), it is possible to automatically derive the cutoff technique that is best suited for the application.




(a) Performance of different recursive computations (Quick-sort and Strassen's algorithm) implemented through the DAC pattern via different back-end runtimes. The plots show the completion time with different parallelism degrees.



(b) Completion time in stencil based denoiser application: random noise (*salt-and-pepper*, left) and gaussian noise (right)



(c) Max input rate sustained by the financial application with various patterns and a different skewness among sub-streams.



(d) FastFlow Pool pattern performances on Intel Sandy Bridge multicore with 24 cores 2-way Hyperthreading.

Fig. 6: Experimental results related to the RePhrase pattern set

5.2 Stencil

In [2] we discussed the FastFlow implementation of a loop-of-stencil-reduce pattern, targeting iterative data parallel computations on heterogeneous multicores. We showed that various iterative kernels can be easily and effectively parallelised by using the Loop-of-stencil-reduce on the available GPUs by ex-

exploiting the OpenCL capabilities of the **FastFlow** parallel framework. We focused on capturing stencil iteration as a pattern, and on its integration in the established **FastFlow** pattern framework. The pattern proved to be quite efficient on modern multicore architectures. Fig. 6b shows the completion times achieved on a 24 core Sandy Bridge machine when executing a video denoiser application using the stencil pattern [4].

5.3 Window-based Streaming Patterns

Data stream processing applications process unbounded data streams coming from a plurality of sensor devices. Input items received at high speed are usually accumulated by updating an internal state of the pattern (e.g., a sliding window containing the most recent data) and by applying a user-defined function periodically, e.g., at each window triggered according to the activation semantics (time-based, count-based or hybrid). When the input stream conveys data items belonging to different logical sub-streams, natural parallelism can be exploited among the computations on windows of different sub-streams.

The **keyedSF** pattern has been adopted in our previous work [11,9] in order to parallelise a high-frequency trading application. The application is fed by a continuous stream of financial ticks that can be *trades*, i.e. closed transactions with a price, a stock symbol and a volume (number of stocks), and *quotes*, that is buy or sell proposals with a proposed price, a stock symbol and a volume. The goal of the application is to automatically discover trading opportunities by analysing the market feeds in near real-time. The computation maintains a sliding window of the most recent data items of each stock symbol (sub-stream) and executes a continuous query at each new window activation. We used count-based windows of 1000 tuples with a refresh slide of 25 new data items. The query computes a least squares curve fitting using the well-known Levenberg-Marquardt algorithm.

From the performance viewpoint the scalability of this parallel pattern depends greatly on the frequency distribution of the sub-streams, because all the windows of the same sub-stream are computed sequentially. Several experiments were performed to evaluate the performance of this pattern under various conditions. Fig. 6c shows the result of an experiment performed on an Intel Ivy Bridge dual-socket multicore workstation featuring 24 cores. The figure depicts the maximum stream rate that the pattern is able to sustain without being a bottleneck by running the application with as many threads as there are cores. We also report the peak performance achieved with a single-threaded implementation of the whole application. While the scalability is almost ideal with a uniform (*Uniform* in Fig. 6c) probability distribution among stock symbols, in more realistic scenarios with a realistic skewness (*Real*) and a heavy skewness (*Skewed*), the scalability of the **keyedSF** pattern drops significantly due to load imbalance.

The figure also shows the performance achieved under the same execution conditions by an alternative implementation based on the **windowedSF** pattern.

In this case the pattern exploits parallelism among windows within the same logical sub-stream by suitably scheduling data items to worker threads in such a way as to execute in parallel consecutive windows with the same stock symbol. The result is a pattern more sophisticated in its data distribution, i.e. an emitter thread is in charge of multicasting each data item to a subset of the workers. However, the performance and load balancing is not affected by the frequency of the sub-streams (the pattern works well also with one stock symbol in the extreme case). This behaviour is evident in the figure, where the peak rate with this second solution is the best under a heavy skewness, while with the uniform distribution and in the real skewness case the `keyedSF` pattern is the winner owing to the more efficient point-to-point distribution of data items to the worker threads.

5.4 Pool

In [1] we designed and implemented implementations of the different variants of the pool pattern in C++/`FastFlow`, as well as in Erlang/`skel` [28]. Both implementations have been used to run experiments on top of state-of-the-art shared memory multicore servers. A full set of experiments has been discussed assessing the features of the pool pattern as well as the efficiency and scalability of the pattern when used to implement various parallel applications. In particular, we have demonstrated that reasonable performances may be achieved with modest programming effort while noting that, in certain cases, manual, ad hoc optimisation of the parallel code taking into account the specific target architecture features may lead to further minor performance improvement. The typical performance figures achieved are exemplified in Fig. 6d. In this case we plot the completion times achieved in the execution of a synthetic benchmark when the number of processing elements (threads) used in the different phases of the pool pattern implementation vary. In particular, the triples (x, y, z) on the x-axis represent the number of threads used in the selection, evolution and filtering phases where the individuals submitted to evolution are selected from the whole population, their evolution is computed and the evolved individuals to be included back in the populations are selected, respectively.

6 Conclusions

We have discussed an extended parallel pattern set designed to support data intensive applications on heterogeneous architectures build of state-of-the-art shared memory multicores and GPUs. We outlined the expressive power of the set, in terms of the range of applications that may be programmed using the patterns and in terms of the programming effort required to implement these applications as compared to the effort involved when using more traditional parallel programming frameworks. Finally, we presented some existing experimental results relative to high level patterns in the pattern set.

Overall, the *RePhrase* pattern set has been demonstrated a) to be sufficient to cover all the use cases adopted in the project as well as the applications included in the well known parallel benchmark suite ParSec, and b) to be able to provide performance comparable to, and in some cases better than, the performance achieved using lower level, *non*-patterned, parallel programming frameworks. This latter point is even more important considering the notable improvement in programmability and the wider range of parallel patterns in the *RePhrase* pattern set as compared to *de facto* standard parallel programming frameworks.

Acknowledgements This work has been partially funded by the EU H2020-ICT-2014-1 project No. 644235 *RePhrase* “Refactoring Parallel Heterogeneous Resource-Aware Applications” (<http://www.rephrase-ict.eu>).

References

1. M. Aldinucci, S. Campa, M. Danelutto, P. Kilpatrick, and M. Torquati. Pool evolution: A parallel pattern for evolutionary and symbolic computing. *Int. J. Parallel Program.*, 44(3):531–551, June 2016.
2. M. Aldinucci, M. Danelutto, M. Drocco, P. Kilpatrick, C. Misale, G. Peretti Pezzi, and M. Torquati. A parallel pattern for iterative stencil + reduce. *Journal of Supercomputing*, pages 1–16, 2016.
3. M. Aldinucci, M. Danelutto, P. Kilpatrick, and M. Torquati. FastFlow: High-Level and Efficient Streaming on Multicore. In S. Pillana and F. Xhafa, editors, *Programming multi-core and many-core computing systems*, Parallel and Distributed Computing, chapter 13. John Wiley & Sons, Inc., 2017.
4. M. Aldinucci, G. Peretti Pezzi, M. Drocco, C. Spampinato, and M. Torquati. Parallel visual data restoration on multi-GPGPUs using stencil-reduce pattern. *Int'l Journal of High Performance Computing Application*, 2015.
5. H. Andrade, B. Gedik, K.-L. Wu, and P. S. Yu. Scale-Up Strategies for Processing High-Rate Data Streams in System S. In *Proceedings of the 2009 IEEE Int'l Conference on Data Engineering*, ICDE '09, pages 1375–1378. IEEE Computer Society, 2009.
6. H. C. Andrade, B. Gedik, and D. S. Turaga. *Fundamentals of Stream Processing: Application Design, Systems, and Analytics*. Cambridge University Press, 2014.
7. K. Asanovic, R. Bodik, J. Demmel, T. Keaveny, K. Keutzer, J. Kubiawicz, 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.
8. B. Babcock, S. Babu, M. Datar, R. Motwani, and J. Widom. Models and issues in data stream systems. In *Proceedings of the twenty-first ACM symposium on Principles of database systems*, PODS '02, pages 1–16, New York, NY, USA, 2002. ACM.
9. M. Danelutto, T. De Matteis, G. Mencagli, and M. Torquati. Data stream processing via code annotations. *The Journal of Supercomputing*, pages 1–15, 2016.
10. M. Danelutto and M. Torquati. Structured parallel programming with “core” fastflow. In V. Zsó�, Z. Horváth, and L. Csató, editors, *Central European Functional Programming School*, volume 8606 of *LNCS*, pages 29–75. Springer, 2015.
11. T. De Matteis and G. Mencagli. Parallel patterns for window-based stateful operators on data streams: An algorithmic skeleton approach. *Int'l Journal of Parallel Programming*, 45(2):382–401, Apr 2017.
12. D. De Sensi, T. De Matteis, M. Torquati, G. Mencagli, and M. Danelutto. Bringing parallel patterns out of the corner: The p3 arsec benchmark suite. *ACM Trans. Archit. Code Optim.*, 14(4):33:1–33:26, Oct. 2017.
13. J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, Jan. 2008.

14. D. del Rio Astorga, M. F. Dolz, L. M. Sánchez, J. G. Blas, and J. D. García. A C++ generic parallel pattern interface for stream processing. In *Algorithms and Architectures for Parallel Processing - 16th Int'l Conference, ICA3PP 2016, Granada, Spain, December 14-16, 2016, Proceedings*, pages 74–87, 2016.
15. A. Duran, E. Ayguadé, R. M. Badia, J. Labarta, L. Martinell, X. Martorell, and J. Planas. Ompps: a proposal for programming heterogeneous multi-core architectures. *Parallel Processing Letters*, 21(2):173–193, 2011.
16. A. Duran, J. Corbalán, and E. Ayguadé. An adaptive cut-off for task parallelism. In *Proceedings of the 2008 ACM/IEEE Conference on Supercomputing, SC '08*, pages 36:1–36:11, Piscataway, NJ, USA, 2008. IEEE Press.
17. K. Emoto and K. Matsuzaki. An Automatic Fusion Mechanism for Variable-Length List Skeletons in SkeTo. *Int'l Journal of Parallel Programming*, 42(4):546–563, 2014.
18. J. Enmyren and C. W. Kessler. Skepu: A multi-backend skeleton programming library for multi-gpu systems. In *Proc. of the Fourth Int'l Workshop on High-level Parallel Programming and Applications, HLPP '10*, pages 5–14. ACM, 2010.
19. S. Ernsting and H. Kuchen. Algorithmic skeletons for multi-core, multi-gpu systems and clusters. *IJHPCN*, 7(2):129–138, 2012.
20. J. Korinth, D. de la Chevallierie, and A. Koch. An open-source tool flow for the composition of reconfigurable hardware thread pool architectures. In *23rd IEEE Annual Int'l Symposium on Field-Programmable Custom Computing Machines, FCCM 2015, Vancouver, Canada, May 2-6, 2015*, pages 195–198. IEEE Computer Society, 2015.
21. D. Laney. 3D data management: Controlling data volume, velocity, and variety. Technical report, META Group, February 2001.
22. N. Marz and J. Warren. *Big Data: Principles and Best Practices of Scalable Realtime Data Systems*. Manning Publications Co., Greenwich, CT, USA, 1st edition, 2015.
23. T. Mattson, B. Sanders, and B. Massingill. *Patterns for Parallel Programming*. Addison-Wesley Professional, first edition, 2004.
24. Microsoft. Task Parallel Library (TPL), 2017. [https://msdn.microsoft.com/en-us/library/dd460717\(v=vs.110\).aspx](https://msdn.microsoft.com/en-us/library/dd460717(v=vs.110).aspx).
25. L. Pappalardo, F. Simini, S. Rinzivillo, D. Pedreschi, F. Giannotti, and A.-L. Barabási. Returners and explorers dichotomy in human mobility. *Nature Communications*, 6:8166+, Sept. 2015.
26. J. Reinders. *Intel Threading Building Blocks*. O'Reilly & Associates, Inc., Sebastopol, CA, USA, first edition, 2007.
27. RePhrase. Report on defined use cases, 2015. RePhrase report D6.3.
28. A Streaming Process-based Skeleton Library for Erlang, 2015. <https://github.com/ParaPhrase/skel>.
29. T. B. Murdoch and A. S. Detsky. The inevitable application of big data to health care. *JAMA*, 309(13):1351–1352, 2013.
30. A. Weiler, M. Grossniklaus, and M. H. Scholl. An Evaluation of the Run-time and Task-based Performance of Event Detection Techniques for Twitter. *Inf. Syst.*, 62(C):207–219, Dec. 2016.
31. T. White. *Hadoop: The Definitive Guide*. O'Reilly Media, Inc., 1st edition, 2009.
32. G. Wilson and R. Irvin. Assessing and comparing the usability of parallel programming systems, 1995. Technical Report, University of Toronto, <http://littlesvr.ca/masters/wp-content/uploads/2010/02/cowichan.pdf>.
33. A. Wright. Big data meets big science. *Commun. ACM*, 57(7):13–15, July 2014.
34. S. Zhao, M. Chandrashekar, Y. Lee, and D. Medhi. Real-time network anomaly detection system using machine learning. In *11th Int'l Conference on the Design of Reliable Communication Networks, DRCN 2015, Kansas City, MO, USA, March 24-27, 2015*, pages 267–270, 2015.