



Challenging the abstraction penalty in parallel patterns libraries

Adding FastFlow support to GrPPI

J. Daniel Garcia¹ · David del Rio¹ · Marco Aldinucci² · Fabio Tordini² · Marco Danelutto³ · Gabriele Mencagli³ · Massimo Torquati³

Published online: 8 April 2019

© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

In the last years, pattern-based programming has been recognized as a good practice for efficiently exploiting parallel hardware resources. Following this approach, multiple libraries have been designed for providing such high-level abstractions to ease the parallel programming. However, those libraries do not share a common interface. To pave the way, GrPPI has been designed for providing an intermediate abstraction layer between application developers and existing parallel programming frameworks like OpenMP, Intel TBB or ISO C++ threads. On the other hand, FastFlow has been adopted as an efficient object-based programming framework that may benefit from being supported as an additional GrPPI backend. However, the object-based approach presents some major challenges to be incorporated under the GrPPI type safe functional programming style. In this paper, we present the integration of FastFlow as a new GrPPI backend to demonstrate that structured parallel programming frameworks perfectly fit the GrPPI design. Additionally, we also demonstrate that GrPPI does not incur in additional overheads for providing its abstraction layer, and we study the programmability in terms of lines of code and cyclomatic complexity. In general, the presented work acts as reciprocal validation of both FastFlow (as an efficient, native structured parallel programming framework) and GrPPI (as an efficient abstraction layer on top of existing parallel programming frameworks).

Keywords Parallel design patterns · Data-intensive computing · Stream computing · Algorithmic skeletons

This work has been partially supported by the European Commission EU H2020-ICT-2014-1 Project RePhrase (No. 644235) and by the Spanish Ministry of Economy and Competitiveness through TIN2016-79637-P “Towards Unification of HPC and Big Data Paradigms”.

✉ J. Daniel Garcia
josedaniel.garcia@uc3m.es

Extended author information available on the last page of the article

1 Introduction

Parallel design patterns have been identified since years as the tool to be leveraged to overcome the difficulties and the scarce productivity typical of parallel programming [4]. Leveraging also on the availability of mature algorithmic skeleton [13]-based programming frameworks [5, 7, 8, 12] this lead to an increasing interest in the possibility of widely using parallel patterns in the development of parallel applications, targeting different kind of architectures, both homogeneous and heterogeneous (e.g., exploiting different kind of accelerators). As an example, different EU funded research projects in the FP7 and H2020 frameworks adopted some kind of algorithmic skeleton/parallel design pattern approach, including FP7 ParaPhrase, Excess [9] and Repara [16] projects, and more recently H2020 RePhrase [17] project.

Existing “structured” parallel programming environments provide the application programmer with high-level tools for the development of parallel applications. In particular, through proper libraries or language extensions they provide ready-to-use programming abstractions modeling common parallel patterns that can be used, alone or in composition, to express the full parallel behavior of a given application. These abstractions encapsulate the more challenging aspects related to parallelism exploitation (synchronization, communication, load balancing, scheduling, etc.) and mainly relieve the application programmers from the burden to manage the target hardware dependent aspects related to the efficient implementation of parallel applications. It is worth pointing out that part of the *state-of-the-art* parallel programming frameworks also include some patterns. OpenMP [15] provides since the very beginning a *map* pattern through its `parallel for pragma`, Intel TBB [20] provides pipeline and farm (master worker) patterns, and Microsoft Parallel Pattern Library [14] also provides different patterns.

However, the lack of (a) a common terminology to denote the different kind of patterns and (b) the lack of a common, recognized and assessed API to use these patterns prevented the diffusion of the concept and a general acknowledgment of the related advantages.

GrPPI (the Generic Parallel Pattern Interface [11]) has been designed as a fully modern C++ compliant interface providing to the application programmers an easy way to parallelize existing code using well-known, efficient stream and data parallel patterns. In its original version, GrPPI targeted C++11 threads to guarantee portability to any ISO C++11 compliant platform. Then, support for OpenMP and TBB has been added. The idea back to GrPPI design was exactly to overcome the two problems just mentioned and, possibly, to provide something that might be eventually included in a future revision of the C++ standard [21]. The inclusion in the language standard would favor the possibility that different parallel programming frameworks providers invest in the implementation of suitable backends with the perspective of getting applications ported for free on their frameworks through a common API.

FastFlow [10] is a structured parallel programming framework-built on top of standard C++ and of the `pthread`s library-providing efficient implementation of both stream and data parallel patterns on top of shared memory architectures. It has been demonstrated to be particularly efficient in the implementation of fine grain computations, due to its extremely optimized communication mechanisms [2]. While simple usage of high-level patterns in **FastFlow** does not require sensibly different expertise than the one required using the same patterns through **GrPPI**, the **FastFlow** API does not reach **GrPPI**'s level of abstraction.

Here, we discuss the integration of **FastFlow** as an additional backend for **GrPPI**. The integration has been challenging being **FastFlow** a framework already providing parallel patterns as objects of a header only library that application programmers and system programmers may use alone or in composition to model a variety of different parallel patterns. The main goal of this paper is to show how the introduction of an extra layer on top of **FastFlow**, namely **GrPPI**, does not necessarily imply a performance penalty and at the same time may simplify the task of expressing parallel applications.

The main contributions of the paper can be summarized as follows:

- We describe the implementation of the **FastFlow** backend for **GrPPI** to demonstrate that **GrPPI** can be used as interface for object-based programming frameworks.
- We analyze the potential performance penalties introduced by the **GrPPI** abstraction layer.
- We evaluate the performance of a set of synthetic benchmarks using **FastFlow** and the different supported **GrPPI** backends.
- We compare the programmability in terms of lines of code and cyclomatic complexity between **GrPPI** and **FastFlow** and analyze the benefits of using both parallel programming frameworks.

The rest of the paper is structured as follows: Sects. 2 and 3 introduce **FastFlow** and **GrPPI**, respectively. Then, Sect. 4 discusses the implementation of the **FastFlow** backend in **GrPPI**. Eventually Sect. 5 shows the experimental results validating the porting and Sect. 6 draws conclusions.

2 FastFlow

FastFlow is a structured parallel programming framework, which is provided as a header only C++ library. **FastFlow** has been designed to target shared memory multi/many core architectures and at the moment being exploits `pthread`s as the main concurrency mechanism, backed up by an original, lock and wait free communication queue that provides ultra efficient inter-thread communications.

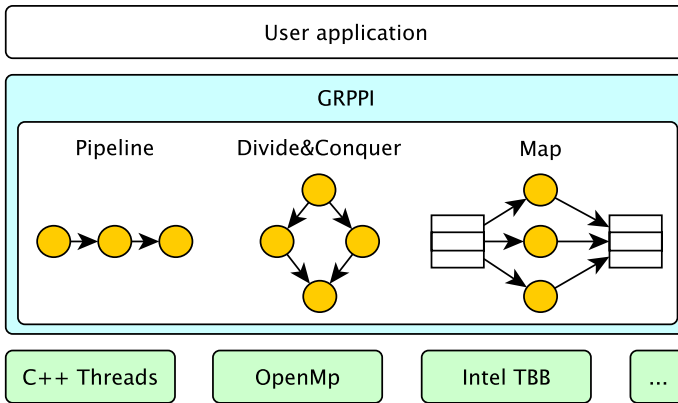


Fig. 1 GrPPI architecture

The design of **FastFlow** is layered:

- A first layer implements the basic communication mechanisms.
- A second layer implements the “core” parallel patterns. These patterns represent notable and efficient parallel patterns exposing a low-level interface suitable to use them both as “application” patterns and as “building block” patterns to be used to provide different patterns to the application programmer.
- A third layer provides “high-level” patterns, with a specific, high level, full C++11 compliant interface, ready to be used within application code.

In either the “core” or the “high-level” pattern layer, **FastFlow** provides (1) stream parallel patterns (pipeline and farm, possibly with feedback); (2) data parallel patterns (parallel for, map, reduce, stencil); and (3) higher level patterns (divide&conquer, evolution pool, windowed stream-processing patterns).

All these patterns are provided as objects the application (or the system) programmer may instantiate and compose to set up a parallel program which is eventually executed by issuing a specific member function call to the top-level pattern. In other words, the declaration of the pattern expression modeling the parallel behavior of the application and the actual execution of the program happen at different places in the program.

Figure 1 shows an excerpt of the code needed to run a pipeline with a farm stage inside in **FastFlow**.

Listing 1: Sample farm+pipeline code in FastFlow

```

1  class Stage1: public ff_node_t<long,long> {
2  public:
3      long* svc(long* task) {
4          // ... here compute res out of task
5          return res;
6      }
7  };
8
9  class Stage2: public ff_node_t<long,long> {...};
10 class Emitter: public ff_node_t<noinput_t, long> {...};
11 class Collector: public ff_node_t<long,nooutput_t> {...};
12
13 int main(int argc, char* argv[]) {
14     // ...
15     // build the pattern expression ...
16     std::vector<std::unique_ptr<ff_node> > workers;
17     for(int i=0;i<nworkers;++i) {
18         // build worker pipeline
19         workers.push_back(make_unique< ff_Pipe<long,long> >(
20             make_unique<Stage1>(), make_unique<Stage2>()));
21     }
22     Emitter E;
23     Collector C;
24     // build the farm with emitter and collector
25     ff_Farm<noinput_t,nooutput_t> farm(std::move(workers), E, C);
26     // Launch farm execution
27     if (farm.run_and_wait_end()<0) {
28         error("running farm\n");
29         return -1;
30     }
31     return 0;
32 }

```

3 Generic and reusable parallel pattern interface

In this section, we introduce GrPPI, a generic and reusable parallel pattern interface for C++ applications [6]. This interface takes full advantage of modern C++ features, generic programming, and meta-programming concepts to act as common interface between different programming models hiding away the complexity behind the use of concurrency mechanisms. Furthermore, the modularity of GrPPI allows to easily integrate new patterns, as well as to combine them to arrange more complex constructions. Thanks to this property, GrPPI can be used to implement a wide range of existing stream-processing and data-intensive applications with relative small efforts, having as a result portable codes that can be executed on multiple platforms. Figure 1 depicts the general view of the GrPPI library. As shown, GrPPI acts as a layer between the user and the different programming models.

Basically, this library provides four main components in order to provide a unified interface for the supported frameworks: (1) type traits, (2) pattern classes, (3) pattern interfaces and (4) execution policies.

Firstly, in order to accomplish the overloading of functions and to allow composition of different patterns, GrPPI provides a set of type traits and `constexpr` functions that are evaluated at compile time. In this sense, these functions can be used

along with the `enable_if` construct from the C++ standard library to conditionally remove functions from compiler resolution. This way, only version of functions meeting specific compile-time requirements on their parameters become available to the compiler. Consequently, this approach allows providing the same interface for different implementations.

On the other hand, GrPPI provides a set of classes that represent each of the supported patterns in order to allow compositions. These objects store references to the necessary functions and information (e.g., concurrency degree) related to the pattern configuration. Thus, they can be composed in order to express complex constructions that cannot be represented by making use of a single pattern.

Focusing on the pattern interfaces, GrPPI provides a set of functions for each supported parallel pattern and offers two different alternatives for pattern execution and composition. Both alternatives take the user functions that will be executed accordingly to the pattern and configuration parameters as function arguments, however, they differ in the first argument. The first version takes an execution policy in order to select the backend which will be used for executing the pattern. On the other hand, the interface designed for pattern composition does not receive the execution policy, as it will use the one from its outer pattern. Instead of launching pattern execution, this version returns the pattern representation that will be used for the composition. Listing 2 shows both interface alternatives for the pipeline pattern.

Listing 2: Pipeline interfaces.

```

1 // Interface for pattern execution
2 template <typename Execution, typename ... Transformers>
3 void pipeline(const Execution & ex, Transformers && ... transform_ops);
4
5 // Interface for pattern composition
6 template <typename ... Transformers>
7 pipeline_t<Transformers...> pipeline(Transformers && ... transform_ops);

```

Finally, a key point of GrPPI is the ability to easily switch between different programming framework implementations for a given pattern. This is achieved by providing a set classes that encapsulate the actual pattern implementations in a specific framework. This way, by only changing the execution policy provided as first argument to the pattern call, the framework used underneath is selected accordingly. The current GrPPI version provides support for sequential, C++ threads, OpenMP and Intel TBB frameworks.

In this paper, we extend the set of GrPPI execution policies in order to support FastFlow as a new backend.

4 FastFlow under GrPPI

In this section, we discuss how the FastFlow framework has been integrated into the GrPPI interface.

FastFlow natively supports high-level stream and data parallel programming patterns built on top of a *lock-free* run-time designed to boost performance on multi-core and many-cores architectures. This makes it perfectly suitable for being used as a backend in the **GrPPI** interface. **FastFlow** provides different high-level patterns as well as different lower level patterns (parallel applications building blocks) that may be both used to implement quite efficient parallel applications. In this section, we will go through the process of implementing **GrPPI**'s interface by means of **FastFlow** patterns, and we will provide an evaluation of the performance of the integration of **FastFlow** into **GrPPI** by comparing the pure **FastFlow** implementation against the wrapped one, in order to identify possible overheads.

4.1 Design

GrPPI provides parallel patterns to support both stream-processing and data-intensive applications. These parallel patterns can be nested and composed together, in order to model more complex behaviors. **GrPPI** currently supports different patterns, including: *data parallel* patterns (map, reduce, mapreduce, stencil), *task-parallel* patterns (divide & conquer), and *stream parallel* patterns (pipeline, farm, stream filter, stream reduction, stream iteration).

The parallel semantics associated to these patterns will not be discussed here. The interested reader may refer to the *RePhrase* deliverables available at <https://rephrase-ict.eu>, in particular in D2.1 [19] and D2.5 [18].

FastFlow natively supports most of the aforementioned skeletons, which are implemented by properly combining instances of the `ff_node` class, which is the main building block upon which every pattern can be designed. In particular, the `ff_node` class implements the abstraction of an independent concurrent activity receiving input tasks from an input queue, processing them and delivering the related results to the output queue. For instance, `ff_pipeline` and `ff_farm` are **FastFlow**'s core patterns built on top of `ff_node`, which represent the most basic and generic stream parallel patterns and *de facto* orchestrate different kind of concurrent activities: stages, in the case of the pipeline pattern, or scheduler, workers and result collector, in the case of the farm pattern.

In its most recent fully C++11 compliant API **FastFlow** introduced enhanced features supporting a better control of the programming patterns typing system. Starting from its building blocks, the `ff_node_t<TypeIn, TypeOut>` becomes a typed abstract container for parallel activities, useful to enforce better type checking on streaming patterns, while among the high-level patterns `ff_Farm<TypeIn, TypeOut>` and `ff_Pipe<TypeIn, TypeOut>` provide typed Farms and Pipelines.

GrPPI takes full advantage of modern C++ features, making an extensive use of template meta-programming and generic programming concepts, that surely help code minimization and leverage automatic compile-time optimizations. The transformations applied to input data are described using C++ lambda functions, where

the callable entities are forwarded via template function parameters, and the body of the lambda may be completely inlined inside a specific specialization of the template function that accepts it. Here, the compiler will optimize the whole body of the function: in performance-critical applications, this solution certainly brings remarkable benefits.

Pattern interfaces are described as function templates, making them more flexible and usable with any data type. The extensive use of variadic templates allows an arbitrary number of data sets to be used by a pattern, and also facilitates handling an arbitrary number of stages/components as typical in a Pipeline pattern, by taking a sequence of callable entities passed as arguments to the function: template overloading guarantees that the right case is matched for every combination of Pipeline stages.

Each pattern function contains an `Execution` type, which represents the backend that will eventually power pattern execution (see Listing 3). For instance, it can be set to operate sequentially, or in parallel by means of one of the supported parallel execution policies. The `FastFlow` backend is fully compliant with these techniques: it has been injected among the supported execution type, and its declaration updated among the type traits and meta-functions that `GrPPI` applies to assert correctness of execution.

Listing 3: GrPPI's Pipeline interface

```

1  template <typename Execution, typename Generator, typename ... Transformers,
2         requires_execution_supported<Execution> = 0>
3  void pipeline(const Execution & ex, Generator && gen_op, Transformers && ... transf_ops)
4  {
5      ex.pipeline(std::forward<Generator>(gen_op), std::forward<Transformers>(transf_ops)...);
6  }

```

The programming framework to be used to support the execution of a high-level `GrPPI` pattern is described specifying the proper `parallel_execution_*` object as a pattern parameter. Each `parallel_execution_*1` class provides member functions for interacting with the underlying parallel framework and contains the actual definition of patterns. Listing 4 reports an excerpt of the `parallel_execution_ff` class, which reflects the way the other backends have actually been designed: by default the concurrency degree is set to the number of available physical cores, but it could be changed via proper access member functions or by using the specific constructor taking the concurrency degree as a parameter.

Listing 4: Parallel execution with FastFlow

```

1  class parallel_execution_ff {
2  public:
3      parallel_execution_ff() noexcept :
4          parallel_execution_ff{static_cast<int>(std::thread::hardware_concurrency())}
5      {}
6      parallel_execution_ff(int degree) noexcept :
7          concurrency_degree_{degree}
8      {}
9
10     // Members for access and modification
11
12     template <typename InputIterator, typename Identity, typename Combiner>
13     auto reduce(InputIterator first, std::size_t sequence_size,
14                Identity && identity, Combiner && combine_op) const;
15
16     // other patterns...
17
18     template <typename Generator, typename ... Transformers>
19     void pipeline(Generator && generate_op, Transformers && ... transform_op) const;
20 }

```

Stream parallel and data parallel patterns are declared and implemented within the execution type: adding FastFlow to GrPPI's data parallel patterns is a rather straightforward job, as they can all be built on top of FastFlow's ParallelFor skeleton. Stream parallel patterns require a slightly bigger effort, as the existing Pipeline pattern has to be used to support the "pattern-matching" mechanism that drives template overloading, leading to the definition of the streaming patterns as composition of Pipelines with a minimum of 3 stages.

Concerning the Divide-and-Conquer (DAC) pattern, FastFlow's DAC pattern does not fully correspond to GrPPI's one; thus, the porting requires some adjustments, due to differences in their programming interface. Further details will be given in the following section.

4.2 Implementation

4.2.1 Data parallel patterns

FastFlow provides the ParallelFor and ParallelForReduce skeletons, designed to efficiently exploit loop parallelism over input data collections. The FastFlow ParallelFor has obtained comparable, or even better performance results with respect to those achieved with well-known frameworks, such as OpenMP or Intel TBB [1, 3], which were already part of GrPPI's backend. As previously mentioned, a FastFlow backend to data parallel patterns is rather straightforward, as they can all be built on top of FastFlow's ParallelFor, even though it requires some caution on pointers arithmetic.

Listing 5 shows the Map pattern built on top of the `ParallelFor`: the object is constructed with the proper concurrency degree, enabling a non-blocking run-time. The grain size also enables a dynamic scheduling of tasks to the threads, and chunks of no more than *grain* iterations at a time are computed by each thread, while the chunk of data is assigned to worker threads dynamically.

Listing 5: Map pattern for FastFlow back-end

```

1  template <typename ... InputIterators, typename OutputIterator, typename Transformer>
2  void parallel_execution_ff::map(std::tuple<InputIterators...> firsts,
3      OutputIterator first_out, std::size_t sequence_size,
4      Transformer transform_op) const
5  {
6      ff::ParallelFor pf{concurrency_degree_, true};
7      long step = 1;
8      long grain = sequence_size/concurrency_degree_;
9
10     pf.parallel_for_idx(0, sequence_size, step, grain,
11         [&](const long start_, const long stop_, const int thid) {
12             for (size_t it_ = start_; it_ < stop_; ++it_)
13                 *(first_out+it_) = apply_iterators_indexed(transform_op, firsts, it_);
14         },
15         concurrency_degree_);
16 }

```

The `ParallelFor` body traverses each chunk with indexes ranges, which are controlled using pointer arithmetic: the `apply_iterators_indexed` is a meta-function provided by GrPPI that applies a callable entity (`transform_op`) to the values obtained from the iterators in a tuple by indexing. Results of the callable are written directly into the output data structure.

Similarly, the Reduce pattern is built on top of the `ParallelFor Reduce`: the reduction is executed in two phases: a partial reduction phase first runs in parallel, while the second phase reduces partial results in series, returning the final output. Iterations are scheduled to worker threads in blocks of at least *grain* iterations.

The Map/Reduce pattern computes a Map-like operation on the input data set, and then applies a Reduce operation on intermediate results.¹ It exploits the intrinsic parallelism provided by the Map and the Reduce patterns, and the implementation is thus a combination of two phases: the Map phase stores its result in a temporary collection, which then undergoes the Reduce phase, and the final result is returned (see Listing 6).

¹ This pattern is not the titled Google's MapReduce, which exploits *key-value* pairs to compute problems that can be parallelized by mapping a function over a given data set or stream of data, and then combining the results.

Listing 6: Map/Reduce pattern for FastFlow back-end

```

1  template <typename ... InputIterators, typename Identity, typename Transformer,
2         typename Combiner>
3  auto parallel_execution_ff::map_reduce(std::tuple<InputIterators...> firsts,
4         std::size_t sequence_size, Identity && identity,
5         Transformer && transform_op, Combiner && combine_op) const
6  {
7      std::vector<Identity> partials(sequence_size);
8      map(firsts, partials.begin(), sequence_size, std::forward<Transformer>(transform_op));
9
10     return reduce(partials.begin(), sequence_size,
11         std::forward<Identity>(identity), std::forward<Combiner>(combine_op));
12 }

```

The Stencil pattern (see Listing 7) is a generalization of the Map pattern, except that it also performs transformations on a set of neighbors in a given coordinate of the input data set. Just like the Map pattern, it is implemented on top of the `ParallelFor`, with the difference that for each element, the result of the transformation is combined with the result of a function applied to the neighboring elements of the current one. The pattern can deal with multiple input data sets, which can be specified by means of variadic parameters in its declaration.

Listing 7: Stencil pattern for FastFlow back-end

```

1  template <typename ... InputIterators, typename OutputIterator,
2         typename StencilTransformer, typename Neighbourhood>
3  void parallel_execution_ff::stencil(std::tuple<InputIterators...> firsts,
4         OutputIterator first_out, std::size_t sequence_size,
5         StencilTransformer&& transform_op, Neighbourhood&& neighbour_op) const
6  {
7      ff::ParallelFor pf(concurrency_degree_, true);
8      long step = 1;
9      long grain = sequence_size/concurrency_degree_;
10
11     pf.parallel_for_idx(0, sequence_size, step, grain,
12         [&](long start, long stop, int thid) {
13         const auto first_it = std::get<0>(firsts);
14         for (size_t i = start; i < stop; ++i) {
15             auto next_chunks = iterators_next(firsts, i);
16             *std::next(first_out, i) = transform_op(std::next(first_it, i),
17                 apply_increment(neighbour_op, next_chunks));
18         }
19     },
20     concurrency_degree_);
21 }

```

4.2.2 Stream parallel patterns

In streaming parallelism, the *end-of-stream* is a special event used to manage termination that has to be orderly notified to the involved concurrent activities, so that they know the computation is terminated and they could start the appropriate closing actions. For this reason, a tag or a marker that identifies the end of data

stream is required. GrPPI uses the `optional<T>` data type, officially introduced in C++17 but already available in C++14 under the *experimental* namespace. This data type is basically a class template that manages an optional contained value, i.e., a value that may or may not be present. The empty `optional` is used to represent the EOS.

On the other hand, *FastFlow* propagates special tags through its actors to notify special events, such as the termination of the data stream, or items to be discarded. For instance, the end-of-stream is identified with the tag `(OutType *) EOS`. This divergence among both interfaces has been handled at the lower level, by wrapping the `ff_node_t` class at the base of patterns' implementation.

A key characteristic of *FastFlow* is the possibility of using their its own memory management mechanism. To allow this, we provide customized versions of operators `new` and `delete` using a special tag `ff_arena`. The use of this internal allocator is highly optimized for *producer/consumer* scenarios, and helps maintaining a low memory footprint during program execution.

By exploiting template overloading, a `node_impl` has been defined for three basic situations in a Pipeline pattern: the first stage generating items, intermediate stages transforming items, and the last stage consuming items. Template `node_impl` is parametrized by input and output data types as well as the callable entity that transforms items.

An intermediate stage is modeled upon the `ff_node_t` class, and models a sequential stage that receives data from the input channel, executes its transformer on the input data and returns its result on the output channel. Both input and output data type need to be specified (see Listing 8).

Listing 8: `node_impl` for intermediate stage

```

1  template <typename Input, typename Output, typename Transformer>
2  class node_impl : public ff::ff_node_t<Input,Output> {
3  public:
4      node_impl(Transformer && transform_op) :
5          transform_op_{transform_op}
6      {}
7
8      Output * svc(Input * p_item) {
9          return new (ff_arena) Output{transform_op_(*p_item)};
10     }
11     private:
12         Transformer transform_op_;
13 };

```

Listing 9 reports the corresponding node for the first stage as a specialization of `node_impl` which does not have an input data type and send items into the data stream as long as its generator returns a value.

Listing 9: node_impl for first stage

```

1  template <typename Output, typename Generator>
2  class node_impl<void,Output,Generator> : public ff::ff_node {
3  public:
4      node_impl(Generator && generate_op) :
5          generate_op_{generate_op}
6      {}
7
8      void * svc(void *) {
9          std::experimental::optional<Output> result{generate_op_()};
10         if (result) { return new (ff_arena) Output{*result}; }
11         else { return EOS; }
12     }
13 private:
14     Generator generate_op_;
15 };

```

The end-of-stream case is adapted by returning the EOS marker when the generator produces an empty object which is GrPPI end-of-stream. Note that this wrapper extends the generic `ff_node` instead of the typed one, because `FastFlow` prohibits to have a node object typed with a `void` input type and a specialized output type: in this case, it falls back to the generic version.

A final stage does not return any data. It is a consumer stage, where input data is stored or printed, and nothing is further placed in the data stream. It is modeled upon the `ff_node_t` class, and no output data type needs to be specified (see Listing 10): if an input value exists, it is consumed and then destroyed (such that memory leaks are avoided), additionally `FastFlow`'s `GO_ON` marker is returned.

Listing 10: node_impl for final stage

```

1  template <typename Input, typename Consumer>
2  class node_impl<Input,void,Consumer> : public ff::ff_node_t<Input,void> {
3  public:
4      node_impl(Consumer && consume_op) :
5          consume_op_{consume_op}
6      {}
7
8      void * svc(Input * p_item) {
9          consume_op_(*p_item);
10         operator delete(p_item, ff_arena);
11         return GO_ON;
12     }
13 private:
14     Consumer consume_op_;
15 };

```

We built the backend for stream parallel patterns using these building blocks. By exploiting template overloading, a Pipeline is constructed out of the stages passed to the variadic parameter in the Pipeline interface (see Listing 3). Except for the very

first stage, which must be the stream generator,² other stages can be any intermediate sequential stage, or another streaming pattern that embeds a parallel logic (e.g., a Farm pattern). For instance, an intermediate stage can be a nested Pipeline as well.

At the base of this mechanism there is FastFlow's Pipeline pattern. By wrapping the `ff_pipeline` class, pattern matching on function templates allows to add stages to the Pipeline in the same order as they have been passed to the variadic parameter in the interface. The only requirement is that each stage must subclass one of the `ff_node` or `ff_node_t` classes, just like the `node_impl` does, as well as FastFlow's Farm pattern. Once the pipeline is constructed, the execution is triggered calling the `run_and_wait_end()` function, which starts the computation and awaits its termination in a cooperative way (see Listing 11).

Listing 11: Construction of the Pipeline

```

1  template <typename Generator, typename ... Transformers>
2  void parallel_execution_ff::pipeline(Generator && generate_op,
3    Transformers && ... transform_ops) const
4  {
5    detail_ff::pipeline_impl pipe{concurrency_degree_, ordering_,
6      std::forward<Generator>(generate_op),
7      std::forward<Transformers>(transform_ops)...};
8    pipe.setFixedSize(false);
9    pipe.run_and_wait_end();
10 }
11
12 class pipeline_impl : public ff::ff_pipeline {
13 public:
14   template <typename Generator, typename ... Transformers>
15   pipeline_impl(int nworkers, bool ordered, Generator && gen,
16     Transformers && ... transform_ops)
17   :
18     nworkers_{nworkers}, ordered_{ordered}, nodes_{}
19   {
20     // Type deduction...
21     auto first_stage = std::make_unique<node_type>(std::forward<Generator>(gen_op));
22     add_node(std::move(first_stage));
23     add_stages<generator_value_type>(std::forward<Transformers>(transform_ops)...);
24   }

```

In order to match the nesting of a stream parallel pattern into the Pipeline, first proper type checking is performed at compile time via template type matching, and then the pattern is constructed and added to the Pipeline. In Listing 11, there are two functions that insert stages into the underlying pipeline object: `add_node` actually adds the stage, while `add_stages` forwards additional stages passed via the variadic parameter, and recurs over the template functions until it finds the right match.

Listing 12 shows the creation of a Farm stage, which is added to the Pipeline before subsequent stages: after type deduction is performed, the actual workers of the Farm are constructed. The workers will be responsible for actually applying Farm's transformation over the input data stream, whose callable is captured at compile time by the `FarmTransformer` parameter. The workers are then added to

² GrPPL assumes all stream parallel computations are pipelines, whose first stage acts as stream generators and the last stage acts as stream absorber. In both cases, the stages may anyway implement some kind of computation on the generated/absorbed stream items.

an ordered.³ Farm (`ff_OFarm`), which also adds default emitter and collector to the Farm. The new pattern is added to the Pipeline, while subsequent stages will be recursively added by matching the proper overloading.

Listing 12: Adding a Farm stage to the Pipeline

```

1  template <typename Input, typename FarmTransformer, template <typename> class Farm,
2         typename ... OtherTransformers,
3         requires Farm<FarmTransformer>> = 0>
4  auto add_stages( Farm<FarmTransformer> && farm_obj,
5                 OtherTransformers && ... other_transform_ops)
6  {
7      std::vector<std::unique_ptr<ff::ff_node>> workers;
8      for(int i=0; i<nworkers_; ++i) {
9          workers.push_back(std::make_unique<worker_type>(
10             std::forward<Farm<FarmTransformer>>(farm_obj)));
11     }
12
13     using node_type = ff::ff_OFarm<Input,output_type>;
14     auto p_farm = std::make_unique<node_type>(std::move(workers));
15     add_node(std::move(p_farm));
16     add_stages<output_type>(std::forward<OtherTransformers>(other_transform_ops)...);
17 }

```

Stream reduce and Stream filter are basically Farm patterns and can easily be nested as Pipeline stages. Due to their semantics, they cannot be implemented as generic Farms, where workers carry on their job and return results, but need to be properly tuned.

The Stream reduce pattern aims at collapsing items appearing on the input stream, and then delivers these results to the output stream. This computation is applied to a subset of the input stream called *window*, where each window is processed independently. In order to accommodate such behavior, a simple Windowed Farm is created, where the emitter is responsible for buffering items coming from the stream, and dispatching them in batch to the workers, as soon as the required window size is reached.

A Windowed Farm may have an *offset* parameter, which may create overlapping regions among two windows. This aspect is also managed by the emitter, that drops items or keeps duplicates, according to what scenario the offset size generates (depending on whether the offset is smaller, equal or greater than the window size). Again, in constructing the Windowed Farm, care has been taken for properly dealing with optional values, which are used to denote the end-of-stream event.

The Stream filter pattern applies a filter over the items in the input stream, such that only those items satisfying the filter pass to the output stream. The filter is computed in parallel using a FastFlow Farm. Despite rather straightforward, an *ad hoc* Farm has been created to reproduce this behavior in FastFlow, to properly deal with presence/absence of output item delivery by farm workers depending on the evaluation of the filter predicate on the corresponding input items.

³ FastFlow provides two different Farm patterns: the *ordered* farm pattern preserves the input/output data stream ordering; the normal farm does not preserve it, for all those computations that, as an example, directly store results in shared memory and the order of the write/updates does not matter.

The Stream iteration pattern applies a transformation to a data item in a loop fashion, until a given condition is satisfied. When the condition is met, the result is sent to the output stream. This pattern has been implemented again using an specialization of `ff_node_t`.

4.2.3 Task-parallel patterns

GrPPI provides a Divide-and-Conquer (DAC) pattern, which applies a multi-branched recursion to break down a problem into several sub-problems, until the base case is reached: at this point distinct sub-problems can be solved in parallel and their solutions merged to form the definitive one. GrPPI comes with two interfaces for calling the DAC pattern, where one of them allows to explicitly pass the Predicate function that leads the divide phase, rather than implicitly use it in the Divide function body.

FastFlow provides a DAC pattern, whose interface requires the Predicate condition for the base case to be explicitly passed as a pure function: this perfectly matches GrPPI; thus, the implementation is rather easy and only requires proper deduction of the output type.

5 Experimental evaluation

In this section, we perform an experimental evaluation of the Map, Reduce, Stencil and Farm in order to compare the patterns' scalability among the different GrPPI backends, the new FastFlow backend and using directly the FastFlow programming framework. Additionally, we have evaluated the programmability in terms of lines of code (LOC) and cyclomatic complexity.

The evaluation has been carried on a machine equipped with two Intel Xeon Ivy Bridge E-2595 with a total of 24 cores running at 2.40 GHz, 03 MB of L3 cache and 128 GB of DDR3 RAM. The OS is a Linux Ubuntu 14.04.2 LST with the kernel 3.13.0-57. As for the C++ compiler we have used the GCC 6.3.0.

5.1 Performance evaluation

To evaluate the performance and programmability, we have designed a set of benchmarks to measure some of the patterns previously described: (1) Map, (2) Reduce, (3) Stencil and (4) Farm. In this section, we present the performance evaluation for the different supported GrPPI backends as well as for direct use of FastFlow. The evaluation allows to analyze the potential overheads introduced by the extra abstraction layer on top of FastFlow. Additionally, we show how GrPPI can be used to compare performance among different backends.

Map benchmark: This benchmark has been used to compute a *daxpy* operation. That is, it computes the operation: $y = \alpha \cdot x + y$, where x and y are vectors of floating point numbers, and α is a random floating point scalar. The result overwrites the




Fig. 2 Execution time of the Map benchmark with respect to the number of cores




Fig. 3 Execution time of the Reduce benchmark with respect to the number of cores




Fig. 4 Execution time of the Stencil benchmark with respect to the number of cores

initial values of vector y . Figure 2 compares execution times using GrPPI against a direct FastFlow implementation. As observed, the execution times scale similarly independently of the framework used underneath.

Reduce benchmark: This benchmark stresses the pattern by computing the addition of a sequence of natural numbers. Figure 3 shows the performance achieved by the different backends with respect to the FastFlow implementation. As can be observed, similarly to the Map pattern benchmark, no significant difference among different backends can be identified. However, in this case, we notice a slight performance degradation when using more than 13 cores. This is due to the inherent overhead related to the data sharing between the two NUMA nodes in which a thread that accesses memory on a remote NUMA node has increased transfer costs.




Fig. 5 Execution time of the Farm benchmark with respect to the number of cores

Stencil benchmark: This benchmark performs a stencil computation over two one-dimensional vectors of the same size (v_1, v_2). The stencil function takes into account the two preceding and the two following neighbors of each item in each vector, and add their values to each item of the first vector. As observed in Fig. 4, the performance is almost the same on all supported backends. Similar to the previous experiment, we notice the NUMA effect when using more than 12 cores in all backends except for the FastFlow implementation. This shows a better handling by FastFlow of NUMA platforms.

Farm benchmark: This benchmark executes some calculation on a stream of natural numbers: for each item in the input stream, workers perform repeated math operations on the input item, and return the result of their computation. As explained above, GrPPI builds all stream parallel patterns as a composition of the Pipeline pattern: for example, a Farm is a three-stages pipeline where the middle stage is a Farm pattern that exploits as many worker threads as the concurrency degree defines. In this comparison, the Farm pattern used with the pure FastFlow benchmark is actually a Pipeline+Farm composition, as reported in Listing 13.

Listing 13: Farm patterns

```

1 // Pure FastFlow version
2 void farm_bench_ff(int cores) {
3   ff::ff_Pipe<> pipe(make_unique<generator>(stream_len),
4     make_unique<ff::ff_OFarm<long>>(worker_func, num_workers), make_unique<Collector>());
5   pipe.run_and_wait_end();
6 }
7
8 // GRPPI version
9 void farm_bench_grppi(const dynamic_execution & e, int cores) {
10  grppi::pipeline(e,
11    [&]() -> experimental::optional<long> {
12      if (idx_in < v.size()) {
13        idx_in++;
14        return v[idx_in-1];
15      }
16      else return {};
17    },
18    grppi::farm(cores, [&](long x) { return do_work(x); }),
19    [&](long x) { output += x; }
20  );
21 }

```

Table 1 Lines of code and cyclomatic complexity of the benchmarks

Benchmark	GrPPI		FastFlow	
	LOC	CCN	LOC	CCN
Map	63	1.4	65	1.4
Reduce	54	1.3	65	1.4
Stencil	76	2.2	73	2.2
Farm	66	1.4	80	1.3

Figure 5 shows the execution time of the Farm benchmark with respect to the number of cores using the different backends. As observed, the TBB backend outperforms the other backends since TBB employs a completely different approach for implementing the pipeline. Intel TBB follows a task-based paradigm with work-stealing. In contrast, the rest of GrPPI backends and FastFlow use dedicated threads as for the generator and collector tasks.

In general, the GrPPI interface does not introduce additional overheads and allows to easily compare the performance by the different backends. This way, GrPPI eases the evaluation and selection of the most adequate backend with a single codebase at the application level and without needing to perform modifications when migrating from one backend to a different one.

5.2 Programmability evaluation

In this section, we study the number of lines of code (LOC) and the average cyclomatic complexity number (CCN) of the GrPPI and FastFlow benchmark implementations. Table 1 shows the number of LOCs and average CCNs as programmability metrics calculated by the `lizard` tool [22]. As observed, since both FastFlow and GrPPI provide high-level interfaces, the number of required LOC and CCN are quite similar. The main difference among both interfaces is that FastFlow is targeted to expert parallel programmers providing more tuning options, while GrPPI provides a more simplified and user-friendly interface aimed to a wider range of application developers, while providing less fine-tuning options. Additionally, the GrPPI implementation supports multiple backends and allows to easily switch among them.

6 Conclusions

In this work, we have presented the design and implementation of the new FastFlow backend for the GrPPI API supporting full pattern-based parallel programming in C++. The experimental evaluation demonstrates that the abstraction layer introduced by the GrPPI interface does not introduce additional overheads and performs equally to using FastFlow directly.

Additionally, we have analyzed the needed programming efforts for implementing a GrPPI application with respect to using an already existing framework. As


we have seen, both approaches provide high-level interfaces that simplify expressing parallel programs. However, the FastFlow interface has been designed targeting expert parallel programmers and provide a way for fine-tuning the application, while GrPPI provides a more functional programming style interface aimed for a wider range of application developers.

References

1. Aldinucci M, Danelutto M, Drocco M, Kilpatrick P, Peretti Pezzi G, Torquati M (2015) The loop-of-stencil-reduce paradigm. In: Proceedings of International Workshop on Reengineering for Parallelism in Heterogeneous Parallel Platforms (RePara). IEEE, Helsinki, Finland, pp 172–177
2. Aldinucci M, Danelutto M, Kilpatrick P, Meneghin M, Torquati M (2012) An efficient unbounded lock-free queue for multi-core systems. In: Euro-Par 2012 Parallel Processing: 18th International Conference, Euro-Par 2012, Rhodes Island, Greece. Springer, New York, pp 662–673
3. Aldinucci M, Peretti Pezzi G, Drocco M, Spampinato C, Torquati M (2015) Parallel visual data restoration on multi-GPGPUs using stencil-reduce pattern. *Int J High Perform Comput Appl* 29:461–472
4. Asanovic K, Bodik R, Demmel J, Keaveny T, Keutzer K, Kubiawicz J, Morgan N, Patterson D, Sen K, Wawrzyniek J, Wessel D, Yelick K (2009) A view of the parallel computing landscape. *Commun ACM* 52(10):56–67
5. Danelutto M, Torquati M (2015) Structured parallel programming with “core” fastflow. In: Zsók V, Horváth Z, Csató L (eds) Central European Functional Programming School, LNCS, vol 8606, Springer, New York, pp 29–75
6. del Rio Astorga D, Dolz MF, Fernández J, García JD (2017) A generic parallel pattern interface for stream and data processing. *Concurr Comput Pract Exp* 29:e4175
7. Ernsting S, Kuchen H (2014) A scalable farm skeleton for hybrid parallel and distributed programming. *Int J Parallel Program* 42(6):968–987
8. Ernstsson A, Li L, Kessler C (2017) Skepu2: flexible and type-safe skeleton programming for heterogeneous parallel systems. *Int J Parallel Program*
9. Excess home page (2017) <http://www.excess-project.eu/>. Accessed Sept 2018
10. FastFlow home page (2017) <http://calvados.di.unipi.it/>. Accessed Sept 2018
11. GrPPI github (2017) <https://github.com/arcosuc3m/grppi>. Accessed Sept 2018
12. Haidi M, Gorlatch S (2018) High-level programming for many-cores using C++14 and the STL. *Int J Parallel Program* 46:23–41
13. Kessler C, Gorlatch S, Enmyren J, Dastgeer U, Steuwer M, Kegel P (2017) Skeleton programming for portable ManyCore computing. In: Programming multicore and manycore computing systems. Wiley, Hoboken
14. Microsoft Parallel Pattern Library home page (2017) <https://msdn.microsoft.com/en-us/library/dd492418.aspx>. Accessed Sept 2018
15. OpenMP home page (2017) <http://www.openmp.org/>. Accessed Sept 2018
16. Repara home page (2017) <http://repara-project.eu/>. Accessed Sept 2018
17. Rephrase home page (2017) <https://rephrase-ict.eu>. Accessed Sept 2018
18. Rephrase Project Technical Report. D2.5 Advanced Pattern Set (2017) <https://rephraseeu.weebly.com/uploads/3/1/0/9/31098995/d2-5.pdf>. Accessed Sept 2018
19. Rephrase Project Technical Report. D2.1. Report on Initial Pattern Set (2017) <https://rephraseeu.weebly.com/uploads/3/1/0/9/31098995/d2-1.pdf>. Accessed Sept 2018
20. TBB home page (2017) <https://www.threadingbuildingblocks.org/>. Accessed Sept 2018
21. Wong M, Garcia JD, Keryell R (2018) Supporting Pipelines in C++. Working Paper P1261R0, ISO/IEC JTC1/SC22/WG21
22. Yin T (2018) Lizard: an cyclomatic complexity analyzer tool online; Accessed 10 Nov 2018

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

J. Daniel Garcia¹  · **David del Rio**¹ · **Marco Aldinucci**² · **Fabio Tordini**² · **Marco Danelutto**³ · **Gabriele Mencagli**³ · **Massimo Torquati**³

David del Rio
david.rio@uc3m.es

Marco Aldinucci
aldinuc@di.unito.it

Fabio Tordini
tordini@di.unito.it

Marco Danelutto
marcod@di.unipi.it

Gabriele Mencagli
mencagli@di.unipi.it

Massimo Torquati
torquati@di.unipi.it

¹ University Carlos III of Madrid, Leganes, Spain

² University of Torino, Turin, Italy

³ University of Pisa, Pisa, Italy