

A Programming Model for Hybrid Workflows: combining Task-based Workflows and Dataflows all-in-one

Cristian Ramon-Cortes^a, Francesc Lordan^a, Jorge Ejarque^a, Rosa M. Badia^a

^a*Barcelona Supercomputing Center (BSC)*

Abstract

In the past years, e-Science applications have evolved from large-scale simulations executed in a single cluster to more complex workflows where these simulations are combined with High-Performance Data Analytics (HPDA). To implement these workflows, developers are currently using different patterns; mainly task-based and dataflow. However, since these patterns are usually managed by separated frameworks, the implementation of these applications requires to combine them; considerably increasing the effort for learning, deploying, and integrating applications in the different frameworks.

This paper tries to reduce this effort by proposing a way to extend task-based management systems to support continuous input and output data to enable the combination of task-based workflows and dataflows (Hybrid Workflows from now on) using a single programming model. Hence, developers can build complex Data Science workflows with different approaches depending on the requirements. To illustrate the capabilities of Hybrid Workflows, we have built a Distributed Stream Library and a fully functional prototype extending COMPSs, a mature, general-purpose, task-based, parallel programming model. The library can be easily integrated with existing task-based frameworks to provide support for dataflows. Also, it provides a homogeneous, generic, and simple representation of object and file streams in both Java and Python; enabling complex workflows to handle any data type without dealing directly with the streaming back-end.

During the evaluation, we introduce four use cases to illustrate the new capabilities of Hybrid Workflows; measuring the performance benefits when processing data continuously as it is generated, when removing synchronisation points, when processing external real-time data, and when combining task-based workflows and dataflows at different levels. The users identifying these patterns in their workflows may use the presented uses cases (and their performance improvements) as a reference to update their code and benefit of the capabilities of Hybrid Workflows. Furthermore, we analyse the scalability in terms of the number of writers and readers and measure the task analysis, task scheduling, and task execution times when using objects or streams.

1. Introduction

For many years, large-scale simulations, High-Performance Data Analytics (HPDA), and simulation workflows have become a must to progress in many scientific areas such as life, health, and earth sciences. In such a context, there is a need to adapt the High-Performance infrastructure and frameworks to support the needs and challenges of workflows combining these technologies [1].

Traditionally, developers have tackled the parallelisation and distributed execution of these applications following two different strategies. On the one hand, task-based workflows orchestrate the execution of several pieces of code (*tasks*) that process and generate data values. These tasks have no state and, during its execution, they are isolated from other tasks; thus, task-based workflows consist of defining the data dependencies among tasks. On the other hand, dataflows assume that tasks are persistent executions with a state that continuously receive/produce data values (streams). Through dataflows, developers describe how the tasks communicate to each other.

Regardless of the workflow type, directed graphs are a useful visualisation and management tool. Figure 1 shows the graph representation of a task-based workflow (left) and its equivalent dataflow (right). The task dependency graph consists of a producer task (coloured in pink) and five consumer tasks (coloured

in blue) that can run in parallel after the producer completes. The dataflow graph also has a producer task (coloured in pink), but one single stateful consumer task (coloured in blue) which processes all the input data sequentially (unless the developer internally parallelises it). Rather than waiting for the completion of the producer task to process all its outputs, the consumer task can process the data as it is generated.




Figure 1: Representation of the same workflow using task-based and dataflow patterns (left and right respectively).

The first contribution of this paper is the proposal of a single hybrid programming model capable of executing task-based workflows and dataflows simultaneously. For that purpose, we extend task-based frameworks to support continuous input and output data, and enable the combination of task-based workflows and dataflows (Hybrid Workflows from now on) using the same programming model. Moreover, it allows developers to build complex Data Science pipelines with different approaches depending on the requirements. The evaluation (Section 6) demonstrates that the use of Hybrid Workflows has significant performance benefits when identifying some patterns in task-based workflows; e.g., when processing data continuously as it is generated, when removing synchronisation points, when processing external real-time data, and when combining task-based workflows and dataflows at different levels. Also, notice that using a single programming model frees the developers from the burden of deploying, using, and communicating different frameworks inside the same workflow.

The second contribution presented in this paper is a Distributed Streaming Library that can be easily integrated with existing task-based frameworks to

provide support for dataflows. The library provides a homogeneous, generic, and simple representation of a stream for enabling complex workflows to handle any kind of data without dealing directly with the streaming back-end. At its current state, the library supports file streams through a custom implementation, and object streams through Kafka [2].

To validate and evaluate our proposal, we have built a prototype that helps us to illustrate the additional capabilities of Hybrid Workflows for complex Data Science applications. This prototype extends COMPSs [3, 4], a mature, general-purpose programming model, and integrates the Distributed Streaming Library.

The rest of the paper is organised as follows. Section 2 presents an overview of the related work, and Section 3 introduces the baseline technology on which our solution is built. Next, Section 4 describes the architecture of the Distributed Stream Library and its integration with COMPSs. Section 5 details four use cases to illustrate the new available features, and Section 6 evaluates our proposal measuring performance improvements at application level and the additional runtime overhead. Finally, Section 7 concludes the paper and gives some guidelines for future work.

2. Related Work

Nowadays, state-of-the-art frameworks typically focus on the execution of either task-based workflows or dataflows. Thus, next subsections provide a general overview of the most relevant frameworks for both task-based workflows and dataflows. Furthermore, since our prototype combines both approaches into a single programming model and allows developers to build Hybrid Workflows without deploying and managing two different frameworks, the last subsection details other solutions and compares them with our proposal.

2.1. Task-based frameworks

Although all the frameworks handle the tasks and data transfers transparently, there are two main approaches to define task-based workflows. On the

one hand, many frameworks force developers to explicitly define the application workflow through a recipe file or a graphical interface. FireWorks [5, 6] defines complex workflows using recipe files in Python, JSON, or YAML. It focuses on high-throughput applications, such as computational chemistry and materials science calculations, and provides support arbitrary computing resources (including queue systems), monitoring through a built-in web interface, failure detection, and dynamic workflow management. Taverna [7, 8] is a suite of tools to design, monitor, and execute scientific workflows. It provides a graphical user interface for the composition of workflows that are written in a Simple Conceptual Unified Flow Language (Scufl) and executed remotely by the Taverna Server to any underlying infrastructure (such as supercomputers, Grids or cloud environments). Similarly, Kepler [9, 10] also provides a graphical user interface to compose workflows by selecting and connecting analytical components and data sources. Furthermore, workflows can be easily stored, reused, and shared across the community. Internally, Kepler’s architecture is actor-oriented to allow different execution models into the same workflow. Also, Galaxy [11, 12] is a web-based platform for data analysis focused on accessibility and reproducibility of workflows across the scientific community. The users define workflows through the web portal and submit their executions to a Galaxy server containing a full repertoire of tools and reference data. In an attempt to increase the interoperability between the different systems and to avoid the duplication of development efforts, Tavaxy [13] integrates Taverna and Galaxy workflows in a single environment; defining an extensible set of re-usable workflow patterns and supporting cloud capabilities. Although Tavaxy allows the composition of workflows using Taverna and Galaxy sub-workflows, the resulting workflow does not support streams nor any dataflow pattern.

On the other hand, other frameworks implicitly build the task dependency graph from the user code. Some opt for defining a new scripting language to manage the workflow. These solutions force the users to learn a new language but make a clear differentiation between the workflow’s management (the script) and the processes or programs to be executed. Nextflow [14, 15] enables

scalable and reproducible workflows using software containers. It provides a fluent DSL to implement and deploy workflows but allows the adaptation of pipelines written in the most common scripting languages. Swift [16, 17] is a parallel scripting language developed in Java and designed to express and coordinate parallel invocations of application programs on distributed and parallel computing platforms. Users only define the main application and the input and output parameters of each program, so that Swift can execute the application in any distributed infrastructure by automatically building the data dependencies.

Other frameworks opt for defining some annotations on top of an already existing language. These solutions avoid the users from learning a new language but merge the workflow annotations and its execution in the same files. Parsl [18] evolves from Swift and provides an intuitive way to build implicit workflows by annotating "apps" in Python codes. In Parsl, the developers annotate Python functions (apps) and Parsl constructs a dynamic, parallel execution graph derived from the implicit linkage between apps based on shared input/output data objects. Parsl then executes apps when dependencies are met. Parsl is resource-independent, that is, the same Parsl script can be executed on a laptop, cluster, cloud, or supercomputer. Dask [19] is a library for parallel computing in Python. Dask follows a task-based approach being able to take into account the data-dependencies between the tasks and exploiting the inherent concurrency. Dask has been designed for computation and interactive data science and integration with Jupyter notebooks. It is built on the dataframe data-structure that offers interfaces to NumPy, Pandas, and Python iterators. Dask supports implicit, simple, task-graphs previously defined by the system (Dask Array or Dask Bag) and, for more complex graphs, the programmer can rely in the **delayed** annotation that supports the asynchronous executions of tasks by building the corresponding task-graph. COMPSs [3, 20, 4] is a task-based programming model for the development of workflows/applications to be executed in distributed programming platforms. The task-dependency graph (or workflow) is generated at execution time and depends on the input data and the dynamic execution of the application. Thus, compared with other workflow

systems that are based on the static drawing of the workflow, COMPSs offers a tool for building dynamic workflows, with all the flexibility and expressivity of the programming language.

2.2. Dataflow frameworks

Stream processing has become an increasingly prevalent solution to process social media and sensor devices data. On the one hand, many frameworks have been created explicitly to face this problem. Apache Flink [21] is streaming dataflow engine to perform stateful computations over data streams (i.e., event-driven applications, streaming pipelines or stream analytics). It provides exactly-once processing, high throughput, automated memory management, and advanced streaming capabilities (such as windowing). Flink users build dataflows that start with one or more input streams (sources), perform arbitrary transformations, and end in one or more outputs (sinks). Apache Samza [22] allows building stateful applications for event processing or real-time analytics. Its differential point is to offer built-in support to process and transform data from many sources, including Apache Kafka, AWS Kinesis, Azure EventHubs, Elasticsearch, and HDFS. Samza users define a stream application that processes messages from a set of input streams, transforms them by chaining multiple operators, and emits the results to output streams or stores. Also, Samza supports at-least-once processing, guaranteeing no data-loss even in case of failures. Apache Storm [23] is a distributed real-time computation system based on the master-worker architecture and used in real-time analytics, online machine learning, continuous computation, and distributed RPC, between others. Storm users define topologies that consume streams of data (spouts) and process those streams in arbitrarily complex ways (bolts), re-partitioning the streams between each stage of the computation however needed. Although Storm natively provides at-least-once processing, it also supports exactly-once processing via its high-level API called Trident. Twitter Heron [24] is a real-time, fault-tolerant stream processing engine. Heron was built as the Storm's successor, meaning that the topology concepts (spouts and bolts) are the same

and has a compatible API with Storm. However, Heron provides better resource isolation, new scheduler features (such as on-demand resources), better throughput, and lower latency.

2.3. Hybrid frameworks

Apache Spark [25] is a general framework for big data processing that was originally designed to overcome the limitations of MapReduce [26]. Among the many built-in modules, Spark Streaming [27] is an extension of the Spark core to evolve from batch processing to continuous processing by emulating streaming via micro-batching. It ingests input data streams from many sources (e.g., Kafka, Flume, Kinesis, ZeroMQ) and divides them into batches that are then processed by the Spark engine; allowing to combine streaming with batch queries. Internally, the continuous stream of data is represented as a sequence of RDDs in a high-level abstraction called Discretized Stream (DStream).

Notice that Spark is based on high-level operators (operators on RDDs) that are internally represented as a DAG; limiting the patterns of the applications. In contrast, our approach is based on sequential programming, which allows the developer to build any kind of application. Furthermore, micro-batching requires a predefined threshold or frequency before any processing occurs; which can be "real-time" enough for many applications, but may lead to failures when micro-batching is simply not fast enough. In contrast, our solution uses a dedicated streaming engine to handle dataflows; relying on streaming technologies rather than micro-batching and ensuring that the data is processed as soon as it is available.

On the other hand, other solutions combine existing frameworks to support Hybrid Workflows. Asterism [28] is a hybrid framework combining `dispel4py` and Pegasus at different levels to run data-intensive stream-based applications across platforms on heterogeneous systems. The main idea is to represent the different parts of a complex application as `dispel4py` workflows which are, then, orchestrated by Pegasus as tasks. While the stream-based execution is managed by `dispel4py`, the data movement between the different execution platforms

and the workflow engine (submit host) is managed by Pegasus. Notice that Asterism can only handle dataflows inside task-based workflows (`dispel4py` workflows represented as Pegasus’ tasks), while our proposal is capable of orchestrating nested task-flows, nested dataflows, dataflows inside task-based workflows, and task-based workflows inside dataflows.

3. Background

To the end of enabling the construction of Hybrid Workflows using the same programming model, we decided to extend an already existing task-based programming model and enable the support for dataflows. Since we have chosen COMPSs as the base workflow manager for our prototype, this section introduces the essential concepts for understanding its programming model and supporting runtime. The second part of this section briefly introduces Kafka since the default backend of the Distributed Streaming Library uses it to support object streams.

3.1. COMPSs

COMP Superscalar (COMPSs) is a programming model based on sequential programming and designed to abstract developers away from the parallelisation and distribution details such as thread creation, synchronisation, data distribution, message passing or fault-tolerance. COMPSs is a task-based model; application developers select a set of methods whose invocations are considered tasks that will run asynchronously in distributed nodes.

As shown in Figure 2, Java is the native programming language to develop COMPSs applications; however, it also provides bindings for Python (Py-COMPSs [29]) and C/C++ [30]. Its programming model is based on annotations that are used to choose class and object methods as tasks. These annotations can be split into two groups:

- **Method Annotations:** Annotations added to the sequential code methods to define them as tasks and potentially execute them in parallel.

- **Parameter Annotations:** Annotations added to the parameters of an annotated method to indicate the direction (IN, OUT, INOUT) of the data used by a task.

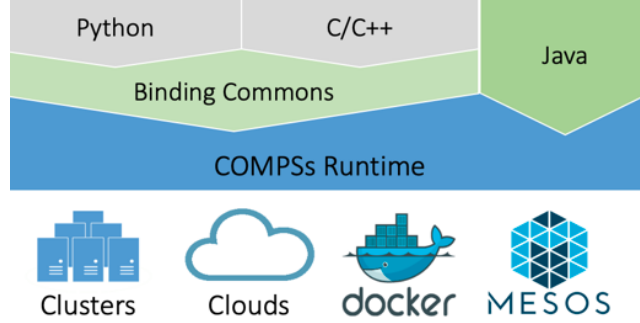


Figure 2: COMPSs overview.

The runtime [20] system supporting the model follows a master-worker architecture. The master node, on which the main code of the application runs, orchestrates the execution of the applications and its tasks on the underlying infrastructure – described in an XML configuration file. For that purpose, it intercepts calls to methods annotated as tasks and, for each detected call, it analyses the data dependencies with previous tasks according to defined parameter annotations. As the result of this analysis, the runtime builds a Directed Acyclic Graph (DAG) where nodes represent tasks and edges represent data dependencies between them; thus, the runtime infers the application parallelism. The master node transparently schedules and submits a task execution on a worker node handling the required data transfers. If a partial failure raises during a task execution, the master node handles it with job re-submission and re-schedule techniques.

COMPSs guarantees the portability across different computing platforms such as clusters, supercomputing machines, clouds, or container-managed infrastructures without modifying the application code. Also, the runtime allows the usage of third-party plugins by implementing a simple interface to extend its usage to new infrastructures or change the scheduling policy.

Furthermore, the COMPSs framework also provides a live-monitoring tool through a built-in web interface. For further details on executions, users can enable the instrumentation of their application using Extrae [31] and generate post-mortem traces that can be analysed with Paraver [32].

3.1.1. Java Syntax

A COMPSs application in Java is composed of three parts:

- **Main application:** Sequential code defining the workflow of the application. It must contain calls to class or object methods annotated as tasks so that, at execution time, they can be asynchronously executed in the available resources.
- **Remote Methods:** Code containing the implementation of the tasks. This code can be in the same file than the application's main code or in one or more separate files.
- **Annotated Interface:** List of annotated methods that can be remotely executed as tasks. It contains one entry per task defining the *Method Annotation*, the object or class method name, and one *Parameter Annotation* per method parameter.

Notice that COMPSs Java applications do not require any special API, pragma or construct since COMPSs instruments the application's code at execution time to detect the tasks defined in the annotated interface. Hence, the COMPSs annotations do not interfere with the applications' code, and all applications can be tested sequentially.

3.1.2. Python Syntax

The Python syntax in COMPSs is supported through a binding: PyCOMPSs. This Python binding is supported by a Binding-commons layer which focuses on enabling the functionalities of the runtime to other languages (currently, Python and C/C++). It has been designed as an API with a set of defined functions.

It is written in C and performs the communication with the runtime through the JNI [33].

In contrast with the Java syntax, all PyCOMPSs annotations are done inline. The Method Annotations are in the form of Python decorators. Hence, the users can add the `@task` decorator on top of a class or object method to indicate that its invocations will become tasks at execution time. Furthermore, the Parameter Annotations are contained inside the Method Annotation.

Listing 1 shows an example of a task annotation. The first line contains the task annotation in the form of a Python decorator while the rest of the code is a regular python function. The parameter `c` is of type `INOUT`, and parameters `a`, and `b` are set to the default type `IN`. The directionality tags are used at execution time to derive the data dependencies between tasks and are applied at an object level, taking into account its references to identify when two tasks access the same object.

```
1  @task(c=INOUT)
2  def multiply(a, b, c):
3      c += a * b
```

Listing 1: PyCOMPSs Task annotation example.

A tiny synchronisation API completes the PyCOMPSs syntax. For instance, as shown in Listing 2, the `compss_wait_on` waits until all the tasks modifying the `result`'s value are finished and brings the value to the node executing the main program (line 4). Once the value is retrieved, the execution of the main program code is resumed. Given that PyCOMPSs is mostly used in distributed environments, synchronising may imply a data transfer from remote storage or memory space to the node executing the main program.

```
1  for block in data:
2      presult = word_count(block)
3      reduce_count(result, presult)
4  final_result = compss_wait_on(result)
```

Listing 2: PyCOMPSs synchronisation API example.

Similarly, the API includes a `compss_open(file_name, mode='r')` to syn-

chronise files, and a `compss_barrier()` to explicitly wait for the completion of all the previous tasks.

3.2. Kafka

Figure 3 illustrates the basic concepts in Kafka and how they relate to each other. *Records* – each blue box in the figure – are key-value pairs containing application-level information registered along with its publication time.

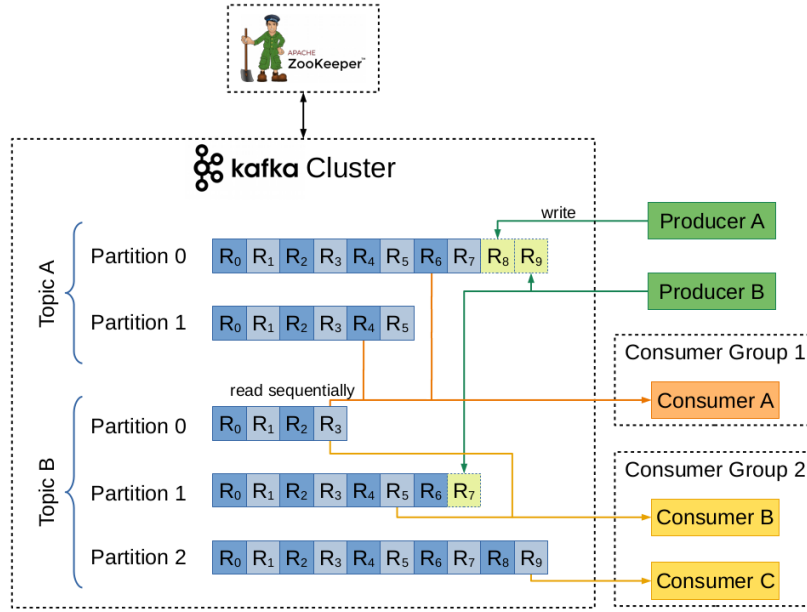


Figure 3: Description of Kafka’s basic concepts.

Kafka users define several categories or *topics* to which records belong. Kafka relies on ZooKeeper [34, 35] to store each topic as a partitioned log with an arbitrary number of partitions and maintains a configurable number of partition replicas across the cluster to provide fault tolerance and record-access parallelism. Each partition contains an immutable, publication-time-ordered sequence of records each uniquely identified by a sequential id number known as the *offset* of the record. The example in the figure defines two topics (*Topic A* and *Topic B*) with 2 and 3 partitions, respectively.

Finally, *Producers* and *Consumers* are third-party application components that interact with Kafka to publish and retrieve data. The former add new records to the topics of their choice, while the latter subscribe to one or more topics for receiving records related to them. Consumers can join in *Consumer groups*. Kafka ensures that each record published to a topic is delivered to at least one consumer instance within each subscribing group; thus, multiple processes on remote machines can share the processing of the records of that topic. Although most often delivered exactly once, records might duplicate when one consumer crashes without a clean shutdown and another consumer within the same group takes over its partitions.

Back to the example in the figure, *Producer A* publishes one record to *Topic A*, and *Producer B* publishes two records, one to *Topic A* and one to *Topic B*. *Consumer A*, with a group of its own, processes all the records in *Topic A* and *Topic B*. Since *Consumer B* and *Consumer C* belong to the same consumer group, they share the processing of all the records from *Topic B*.

Besides the Consumer and Producer API, Kafka also provides the Stream Processor and Connector APIs. The former, usually used in the intermediate steps of the fluent stream processing, allows application components to consume an input stream from one or more topics and produce an output stream to one or more topics. The latter is used for connecting producers and consumers to already existing applications or data systems. For instance, a connector to a database might capture every change to a table.

4. Architecture

Figure 4 depicts a general overview of the proposed solution. When executing regular task-based workflows, the application written following the programming model interacts with the runtime to spawn the remote execution of tasks and retrieve the desired results. Our proposal includes a representation of a stream (*DistroStream* Interface) that provides applications with homogeneous stream accesses regardless of the stream backend supporting them. Moreover,

we extend the programming model and runtime to provide task annotations and scheduling capabilities for streams.

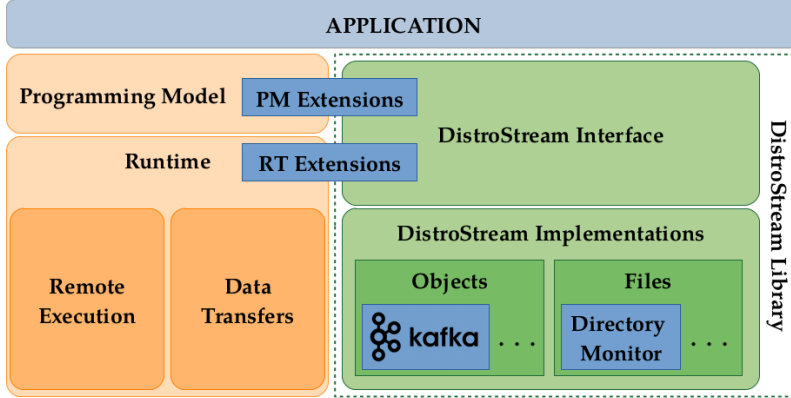


Figure 4: General architecture.

The following subsections discuss the architecture of the proposed solution in a bottom-up approach, starting from the representation of a stream (*DistroStream* API) and its implementations. Next, we describe the *Distributed Stream Library* and its internal components. Finally, we detail the integration of this library with the programming model (COMPSs) and the necessary extensions of its runtime system.

4.1. Distributed Stream Interface

The Distributed Stream is a representation of a stream used by applications to publish and receive data values. Its interface provides a common API to guarantee homogeneity on all interactions with streams.

As shown in Listing 3, the `DistroStream` interface provides a `publish` method for submitting a single message or a list of messages (lines 5 and 6) and a `poll` method to retrieve all the currently available unread messages (lines 9 and 10). Notice that the latter has an optional `timeout` parameter (in milliseconds) to wait until an element becomes available or the specified time expires. Moreover, the streams can be created with an optional `alias` parameter (line 2) to allow different applications to communicate through them. Also, the inter-

```

1  // INSTANTIATION
2  public DistroStream(String alias) throws RegistrationException;
3
4  // PUBLISH METHODS
5  public abstract void publish(T message) throws BackendException;
6  public abstract void publish(List<T> messages) throws BackendException;
7
8  // POLL METHODS
9  public abstract List<T> poll() throws BackendException;
10 public abstract List<T> poll(long timeout) throws BackendException;
11
12 // METADATA METHODS
13 public StreamType getStreamType();
14 public String getAlias();
15 public String getId();
16
17 // STREAM STATUS
18 public boolean isClosed();
19
20 // CLOSE STREAM
21 public final void close();

```

Listing 3: Distributed Stream Interface in Java.

face provides other methods to query stream metadata; such as the stream type (line 13), id (line 14), or alias (line 15). Finally, the interface includes methods to check the status of a stream (line 18), and to close it (line 21).

Due to space constraints, Listing 3 only shows the Java interface, but our prototype also provides the equivalent interface in Python.

4.2. Distributed Stream Implementations

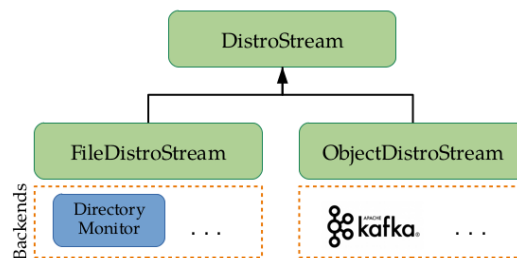


Figure 5: DistroStream class relationship.

As shown in Figure 5, two different implementations of the `DistroStream` API provide the specific logic to support object and file streams. Object streams are suitable when sharing data within the same language or framework. On the

other hand, file streams allow different frameworks and languages to share data. For instance, the files generated by an MPI simulation in C or Fortran can be received through an stream and processed in a Python or Java application.

4.2.1. Object streams

`ObjectDistroStream` (ODS) implements the generic `DistroStream` interface to support object streams. Each ODS has an associated `ODSPublisher` and `ODSConsumer` that interact appropriately with the software handling the message transmission (streaming backend). The ODS instantiates them upon the first invocation of a `publish` or a `poll` method respectively. This behaviour guarantees that the same object stream has different publisher and consumer instances when accessed from different processes, and that the producer and consumer instances are only registered when required, avoiding unneeded registrations on the streaming backend.

At its current state, the available implementation is backed by Kafka, but the design is prepared to support many backends. Notice that the ODS, the `ODSPublisher`, and the `ODSConsumer` are just abstractions to hide the interaction with the underlying backend. Hence, any other backend (such as an MQTT broker) can be supported without any modification at the workflow level by implementing the functionalities defined in these abstractions.

Considering the Kafka concepts introduced in Section 3.2, each ODS becomes a Kafka topic named after the stream id. When created, the `ODSPublisher` instantiates a `KafkaProducer` whose `publish` method builds a new `ProducerRecord` and submits it to the corresponding topic via the `KafkaProducer.send` method. If the `publish` invocation sends several messages, the `ODSPublisher` iteratively performs the publishing process for each message so that Kafka registers it as separated records.

Likewise, a new `KafkaConsumer` is instantiated along with an `ODSConsumer`. Then, the `KafkaConsumer` is registered to a consumer group shared by all the consumers of the same application to avoid replicated messages, and subscribed to the topic named after the id of the stream. Hence, the `poll` method retrieves

a list of `ConsumerRecords` and deserialises their values. To ensure that records are processed exactly-once, consumers also interact with Kafka's `AdminClient` to delete all the processed records from the database.

```
1  // PRODUCER
2  void produce(List<T> objs) {
3      // Create stream (alias is not mandatory)
4      String alias = "myStream";
5      ObjectDistroStream<T> ods = new ObjectDistroStream<>(alias);
6      // Metadata getters
7      System.out.println("Stream Id: " + ods.getId());
8      System.out.println("Stream Alias: " + ods.getAlias());
9      System.out.println("Stream Type: " + ods.getStreamType());
10     // Publish (single element or list)
11     for (T obj : objs) {
12         ods.publish(obj);
13     }
14     ods.publish(objs);
15     // Close stream
16     ods.close()
17 }
18
19 // CONSUMER
20 void consume(ObjectDistroStream<T> ods) {
21     // Poll current elements (without timeout)
22     if (!ods.isClosed()) {
23         List<T> newElems = ods.poll();
24     }
25     // Poll until stream is closed (with timeout)
26     while (!ods.isClosed()) {
27         List<T> newElems = ods.poll(5)
28     }
29 }
```

Listing 4: Object Streams (ODS) example in Java.

Listing 4 shows an example using object streams in Java. Notice that the stream creation (line 5) forces all the stream objects to be of the same type `T`. Internally, the stream serialises and deserialises the objects so that the application can publish and poll elements of type `T` directly to/from the stream. As previously explained, the example also shows the usage of the `publish` method for a single element (line 12) or a list of elements (line 14), the `poll` method with the optional `timeout` parameter (lines 23 and 27, respectively), and the common API calls to close the stream (line 16), check its status (line 22), and retrieve metadata information (lines 7 to 9).

Due to space constraints, the example only shows the ODS usage in Java,

but our prototype provides an equivalent implementation in Python.

4.2.2. File streams

The `FileDistroStream` implementation (FDS) backs the `DistroStream` up to support the streaming of files. Like ODS, its design allows using different backends; however, at its current state, it uses a custom implementation that monitors the creation of files inside a given directory. The *Directory Monitor* backend sends the file locations through the stream and relies on a distributed file system to share the file content. Thus, the monitored directory must be available to every client on the same path.

```
1  // PRODUCER
2  void produce(String baseDir, List<String> fileNames) throws IOException {
3      // Create stream (alias is not mandatory)
4      String alias = "myStream";
5      FileDistroStream<T> fds = new FileDistroStream<>(alias, baseDir);
6      // Publish files (no need to explicitly call the publish
7      // method, the baseDir directory is automatically monitored)
8      for (String fileName : fileNames) {
9          String filePath = baseDir + fileName;
10         try (BufferedWriter writer = new BufferedWriter(new FileWriter(filePath))) {
11             writer.write(...);
12         }
13     }
14     // Close stream
15     fds.close()
16 }
17
18 // CONSUMER
19 void consume(FileDistroStream<T> fds) {
20     // Poll current elements (without timeout)
21     if (!fds.isClosed()) {
22         List<String> newFiles = fds.poll();
23     }
24     // Poll until stream is closed (with timeout)
25     while (!fds.isClosed()) {
26         List<String> newFiles = fds.poll(5)
27     }
28 }
```

Listing 5: File Streams (FDS) example in Java.

Listing 5 shows an example using file streams in Java. Notice that the FDS instantiation (line 5 in the listing) requires a base directory to monitor the creation of files and that it optionally accepts an alias argument to retrieve the content of an already existing stream. Also, files are not explicitly published on

the stream since the base directory is automatically monitored (lines 8 to 13). Instead, regular methods to write files are used. However, the consumer must explicitly call the `poll` method to retrieve a list of the newly available file paths in the stream (lines 22 and 26). As with ODS, applications can also use the common API calls to close the stream (line 15), check its status (lines 21 and 25), and retrieve metadata information.

Due to space constraints, the example only shows the FDS usage in Java, but our prototype provides an equivalent implementation in Python.

4.3. Distributed Stream Library

The Distributed Stream Library (`DistroStreamLib`) handles the stream objects and provides three major components. First, the `DistroStream` API and implementations described in the previous sections.

Second, the library provides the `DistroStream Client` that must be available for each application process. The client is used to forward any stream metadata request to the `DistroStream Server` or any stream data access to the suitable stream backend (i.e., *Directory Monitor*, or *Kafka*). To avoid repeated queries to the server, the client stores the retrieved metadata in a cache-like fashion. Either the Server or the backend can invalidate the cached values.

Third, the library provides the *DistroStream Server* process that is unique for all the applications sharing the stream set. The server maintains a registry of active streams, consumers, and producers with the purpose of coordinating any stream data or metadata access. Among other responsibilities, it is in charge of assigning unique ids to new streams, checking the access permissions of producers and consumers when requesting `publish` and `poll` operations, and notifying all registered consumers when the stream has been completely closed and there are no producers remaining.

Figure 6 contains a sequence diagram that illustrates the interaction of the different Distributed Stream Library components when serving a user petition. The *DistroStream* implementation used by the applications always forwards the requests to the *DistroStream Client* available on the process. The client

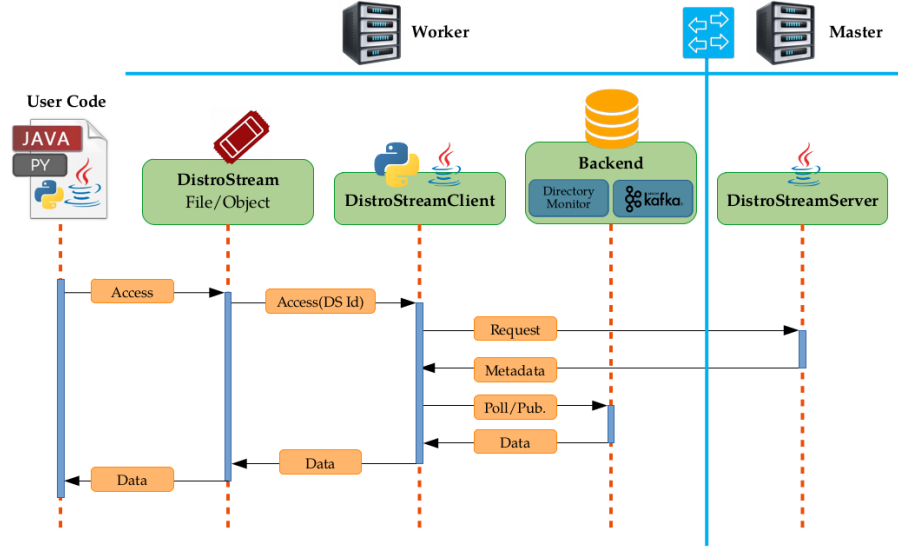


Figure 6: Sequence diagram of the Distributed Stream Library components.

communicates with the *DistroStream Server* for control purposes, and retrieves from the backend the real data.

4.4. Programming Model Extensions

As already mentioned in Section 3.1, the prototype to evaluate Hybrid Workflows is based on the COMPSs workflow manager. At programming-model level, we have extended the COMPSs Parameter Annotation to include a new *STREAM* type.

As shown in Listing 6, on the one hand, the users declare producer tasks (methods that write data into a stream) by adding a parameter of type *STREAM* and direction *OUT* (lines 3 to 7 in the listing). On the other hand, the users declare consumer tasks (methods that read data from a stream) by adding a parameter of type *STREAM* and direction *IN* (lines 9 to 13 in the listing). In the current design, we have not considered *INOUT* streams because we do not imagine a use case where the same method writes data into its own stream. However, it can be easily extended to support such behaviour when required.

```

1  public interface Itf {
2
3      @Method(declaringClass = "Producer")
4      Integer sendMessages(
5          @Parameter(type = Type.STREAM, direction = Direction.OUT)
6          DistroStream stream
7      );
8
9      @Method(declaringClass = "Consumer")
10     Result receiveMessages(
11         @Parameter(type = Type.STREAM, direction = Direction.IN)
12         DistroStream stream
13     );
14 }

```

Listing 6: Stream parameter annotation example in Java.

Furthermore, we want to highlight that this new annotation allows integrating streams smoothly with any other previous annotation. For instance, Listing 7 shows a single producer task that uses two parameters: a stream parameter typical of dataflows (lines 5 and 6) and a file parameter typical of task-based workflows (lines 7 and 8).

```

1  public interface Itf {
2
3      @Method(declaringClass = "Producer")
4      Integer sendMessages(
5          @Parameter(type = Type.STREAM, direction = Direction.OUT)
6          DistroStream stream,
7          @Parameter(type = Type.FILE, direction = Direction.IN)
8          String file
9      );
10 }

```

Listing 7: Example combining stream and file parameters in Java.

4.5. Runtime extensions

As depicted in Figure 7, COMPSs registers the different tasks from the application’s main code through the Task Analyser component. Then, it builds a task graph based on the data dependencies and submits it to the *Task Dispatcher*. The Task Dispatcher interacts with the Task Scheduler to schedule the data-free tasks when possible and, eventually, submit them to execution. The

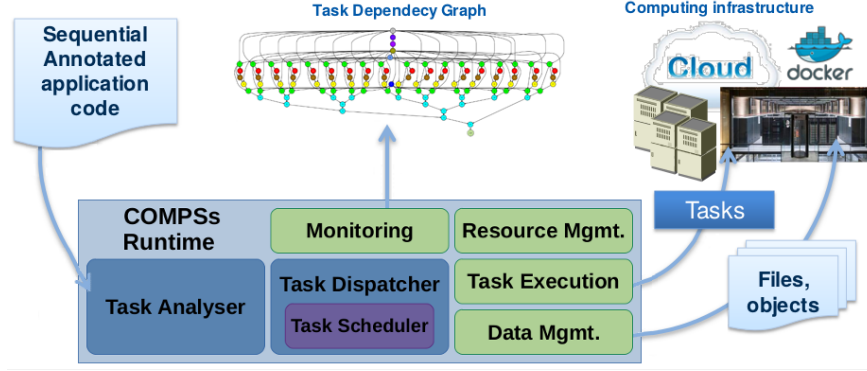


Figure 7: Structure of the internal COMPSs components.

execution step includes the job creation, the transfer of the input data, the job transfer to the selected resource, the real task execution on the worker, and the output retrieval from the worker back to the master. If any of these steps fail, COMPSs provides fault-tolerant mechanisms for partial failures. Also, once the task has finished, COMPSs stores the monitoring data of the task, synchronises any data required by the application, releases the data-dependent tasks so that they can be scheduled, and deletes the task.

Therefore, the new Stream annotation has forced modifications in the Task Analyser and Task Scheduler components. More specifically, notice that a stream parameter does not define a traditional data dependency between a producer and a consumer task since both tasks can run at the same time. However, there is some information that must be stored so that the Task Scheduler can correctly handle the available resources and the data locality. In this sense, when using the same stream object, our prototype prioritises producer tasks over consumer tasks to avoid wasting resources when a consumer task is waiting for data to be produced by a non-running producer task. Moreover, the Task Scheduler assumes that the resources that are running (or have run) producer tasks are the data locations for the stream. This information is used to schedule the consumer tasks accordingly and minimise as much as possible the data transfers between nodes.




Figure 8: COMPSs and Distributed Stream Library deployment.

Regarding the components' deployment, as shown in Figure 8, the COMPSs master spawns the *DistroStream Server* and the required backend. Furthermore, it includes a *DistroStream Client* to handle the stream accesses and requests performed on the application's main code. On the other hand, the COMPSs workers only spawn a *DistroStream Client* to handle the stream accesses and requests performed on the tasks. Notice that the COMPSs master-worker communication is done through NIO [36], while the *DistroStream Server-Client* communication is done through Sockets.

5. Use Cases

Enabling Hybrid task-based workflows and dataflows into a single programming model allows users to define new types of complex workflows. We introduce four patterns that appear in real-world applications so that the users identifying these patterns in their workflows can benefit from the new capabilities and performance improvements of Hybrid Workflows. Next subsections provide in-depth analysis of each use case.

5.1. Use case 1: Continuous data generation

One of the main drawbacks of task-based workflows is waiting for task completion to process its results. Often, the output data is generated continuously

during the task execution rather than at the end. Hence, enabling data streams allows users to process the data as it is generated.

```

1  @constraint(computing_units=CORES_SIMULATION)
2  @task(varargs_type=FILE_OUT)
3  def simulation(num_files, *args):
4      ...
5
6  @constraint(computing_units=CORES_PROCESS)
7  @task(input_file=FILE_IN, output_image=FILE_OUT)
8  def process_sim_file(input_file, output_image):
9      ...
10
11 @constraint(computing_units=CORES_MERGE)
12 @task(output_gif=FILE_OUT, varargs_type=FILE_IN)
13 def merge_reduce(output_gif, *args):
14     ...
15
16 def main():
17     # Parse arguments
18     num_sims, num_files, sim_files, output_files, output_gifs = ...
19     # Launch simulations
20     for i in range(num_sims):
21         simulation(num_files, *sim_files[i])
22     # Process generated files
23     for i in range(num_sims):
24         for j in range(num_files):
25             process_sim_file(sim_files[i][j], output_images[i][j])
26     # Launch merge phase
27     for i in range(num_sims):
28         merge_reduce(output_gifs[i], *output_images[i])
29     # Synchronize files
30     for i in range(num_sims):
31         output_gifs[i] = compss_wait_on_file(output_gifs[i])

```

Listing 8: Simulations' application in Python without streams.

For instance, Listing 8 shows the code of a pure task-based application that launches `num_sims` simulations (line 21). Each simulation produces output files at different time steps of the simulation (i.e., an output file every iteration of the simulation). The results of these simulations are processed separately by the `process_sim_file` task (line 25) and merged to a single GIF per simulation (line 28). The example code also includes the task definitions (lines 1 to 13) and the synchronisation API calls to retrieve the results (line 31).

Figure 9 shows the task graph generated by the previous code when running with 2 simulations (`num_sims`) and 5 files per simulation (`num_files`). The `simulation` tasks are shown in blue, the `process_sim_file` in white and red,

and the `merge_reduce` in pink. Notice that the simulations and the processing of the files cannot run in parallel since the task-based workflow forces the completion of the simulation tasks to begin any of the processing tasks.




Figure 9: Task graph of the simulation application without streaming.

On the other hand, Listing 9 shows the code of the same application using streams to retrieve the data from the simulations as it is generated and forwarding it to its processing tasks. The application initialises the streams (lines 20 to 22), launches `num_sims` simulations (line 25), spawns a process task for each received element in each stream (line 34), merges all the output files into a single GIF per simulation (line 37), and synchronises the final results. The `process_sim_file` and `merge_reduce` task definitions are identical to the previous example. Conversely, the `simulation` task definition uses the `STREAM_OUT` annotation to indicate that one of the parameters is a stream where the task is going to publish data. Also, although the simulation, merge, and synchronisation phases are very similar to the pure task-based workflow, the processing phase is completely different (lines 27 to 34). When using streams, the main code needs to check the stream status, retrieve its published elements, and spawn a `process_sim_task` per element. However, the complexity of the code does not increase significantly when adding streams to an existing application.

Figure 10 shows the task graph generated by the previous code when running with the same parameters than the pure task-based example (2 simulations and 5 files per simulation). The colour code is also the same than the previous

```

1  @constraint(computing_units=CORES_SIMULATION)
2  @task(fds=STREAM_OUT)
3  def simulation(fds, num_files):
4      ...
5
6  @constraint(computing_units=CORES_PROCESS)
7  @task(input_file=FILE_IN, output_image=FILE_OUT)
8  def process_sim_file(input_file, output_image):
9      ...
10
11 @constraint(computing_units=CORES_MERGE)
12 @task(output_gif=FILE_OUT, varargs_type=FILE_IN)
13 def merge_reduce(output_gif, *args):
14     ...
15
16 def main():
17     # Parse arguments
18     num_sims, num_files, output_images, output_gifs = ...
19     # Initialise streams
20     input_streams = [None for _ in range(num_sims)]
21     for i in range(num_sims):
22         input_streams[i] = FileDistroStream(base_dir=stream_dir)
23     # Launch simulations
24     for i in range(num_sims):
25         simulation(input_streams[i], num_files)
26     # Process generated files
27     for i in range(app_args.num_simulations):
28         while not input_streams[i].is_closed():
29             # Process new files
30             new_files = input_streams[i].poll()
31             for input_file in new_files:
32                 output_image = input_file + ".out"
33                 output_images[i].append(output_image)
34                 process_sim_file(input_file, output_image)
35     # Launch merge phase
36     for i in range(app_args.num_simulations):
37         merge_reduce(output_gifs[i], *output_images[i])
38     # Synchronise files
39     for i in range(app_args.num_simulations):
40         output_gifs[i] = compss_wait_on_file(output_gifs[i])

```

Listing 9: Simulations' application in Python with streams.

example: the `simulation` tasks are shown in blue, the `process_sim_file` in white and red, and the `merge_reduce` in pink. Notice that streams enable the execution of the processing tasks while the simulations are still running; potentially reducing the total execution time and increasing the resources utilisation (see Section 6.2 for further details).

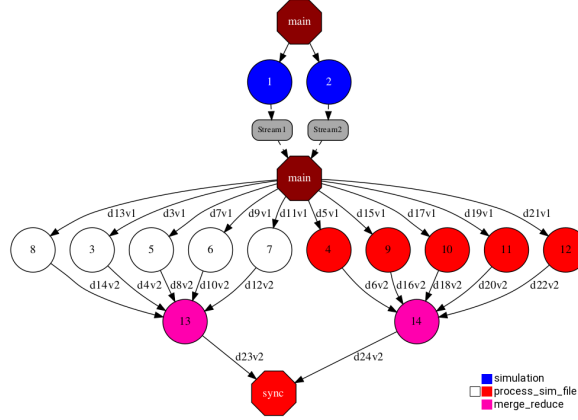


Figure 10: Task graph of the simulation application with streaming.

5.2. Use case 2: Asynchronous data exchange

Streams can also be used to communicate data between tasks without waiting for the tasks' completion. This technique can be useful when performing parameter sweep, cross-validation, or running the same algorithm with different initial points.

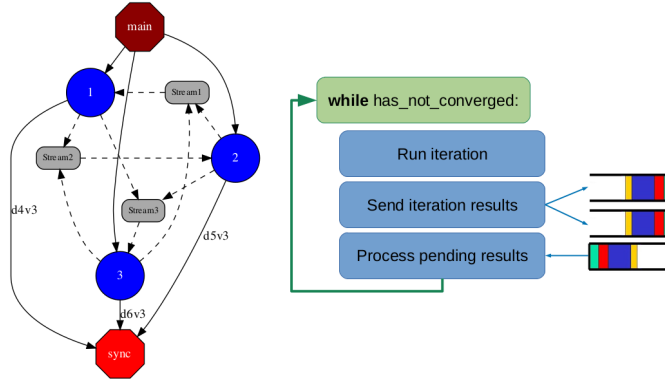


Figure 11: Task graph of the multi-simulations application.

For instance, Figure 11 shows three algorithms running simultaneously that exchange control data at the end of every iteration. Notice that the data exchange at the end of each iteration can be done synchronously by stopping all the simulations, or asynchronously by sending the updated results and processing

the pending messages in the stream (even though some messages of the actual iteration might be received in the next iteration). Furthermore, each algorithm can run a complete task-based workflow to perform the iteration calculus obtaining a nested task-based workflow inside a pure dataflow.

5.3. Use case 3: External streams

Many applications receive its data continuously from external streams (i.e., IoT sensors) that are not part of the application itself. Moreover, depending on the workload, the stream data can be produced by a single task and consumed by many tasks (one to many), produced by many tasks and consumed by a single task (many to one), or produced by many tasks and consumed by many tasks (many to many). The *Distributed Stream Library* supports all three scenarios transparently, and allows to configure the consumer mode to process the data at least once, at most once, or exactly once when using many consumers.




Figure 12: Task graph of the sensor application.

Figure 12 shows an external sensor (*Stream 1* in the figure) producing data

in pink) produces the data, task 2 (in white) filters it, task 3 (in blue) extracts and collects the data, and task 4 (in red) runs a big computation.

Notice that, in the previous use case, the application always has 4 filter tasks. However, in this scenario, the filter task has a nested task-based workflow that accumulates the received data into batches and spawns a new filter task per batch. This technique dynamically adapts the resource usage to the amount of data received by the input stream. Likewise, the big computation task also contains a nested task-based workflow. This shows that users can parallelise some computations internally without modifying the original dataflow.

6. Evaluation

This section evaluates the performance of the new features enabled by our prototype when using data streams against their equivalent implementations using task-based workflows. Furthermore, we analyse the stream writer and reader processes' scalability and load balancing. Finally, we provide an in-depth analysis of the COMPSs runtime performance by comparing the task analysis, task scheduling, and task execution times when using pure task-based workflows or streams.


6.1. Experimental Setup

The results presented in this section have been obtained using the MareNostrum 4 Supercomputer [37] located at the Barcelona Supercomputing Center (BSC). Its current peak performance is 11.15 Petaflops. The supercomputer is composed by 3456 nodes, each of them with two Intel®Xeon Platinum 8160 (24 cores at 2,1 GHz each). It has 384.75 TB of main memory, 100Gb Intel®Omni-Path Full-Fat Tree Interconnection, and 14 PB of shared disk storage managed by the Global Parallel File System.


Regarding the software, we have used DistroStream Library (available at [38]), COMPSs version 2.5.rc1909 (available at [39]), and Kafka version 2.3.0 (available at [40]). We have also used Java OpenJDK 8 131, Python 2.7.13, GCC 7.2.0, and Boost 1.64.0.

6.2. Gain of processing data continuously

As explained in the first use case in Section 5.1, one of the significant advantages when using data streams is to process data continuously as it is generated. For that purpose, Figure 14 compares the Paraver [32] traces of the original COMPSs execution (pure task-based workflow) and the execution using Hybrid Workflows. Each trace shows the available threads in the vertical axis and the execution time in the horizontal axis - 36s in both traces. Also, each colour represents the execution of a task type; corresponding to the colours shown in the task graphs of the first use case (see Section 5.1). The green flags indicate when a simulation has generated all its output files and has closed its associated writing stream. Both implementations are written in Python, and the *Directory Monitor* is set as stream backend.



(a) Original COMPSs execution



(b) Execution with Hybrid Workflows

Figure 14: Paraver traces to illustrate the gain of processing data continuously.

In contrast to the original COMPSs execution (top), the execution using Hybrid Workflows (bottom) executes the processing tasks (white and red) while the simulations (blue) are still running; significantly reducing the total execution time and increasing the resources utilisation. Moreover, the `merge_reduce` tasks (pink) are able to begin its execution even before the simulation tasks are finished, since all the streams have been closed and the `process_sim_file` tasks have already finished.

In general terms, the gain of the implementation using Hybrid Workflows with respect to the original COMPSs implementation (calculated following Equa-

tion 1) is proportional to the number of tasks that can be executed in parallel while the simulation is active. Therefore, we perform an in-depth analysis of the trade-off between the generation and process times. It is worth mentioning that we define the *generation time* as the time elapsed between the generation of two elements of the simulation. Hence, the total duration of the simulation is the generation time multiplied by the number of generated elements. Also, the *process time* is defined as the time to process a single element (that is, the duration of the `process_sim_file` task).

$$Gain = \frac{ExecutionTime_{original} - ExecutionTime_{hybrid}}{ExecutionTime_{original}} \quad (1)$$

The experiment uses 2 nodes of 48 cores each. Since the COMPSs master reserves 12 cores, there are two available workers with 36 and 48 cores respectively. The simulation is configured to use 48 cores, leaving 36 available cores while it is active and 84 available cores when it is over. Also, the process tasks are configured to use one single core.




Figure 15: Average execution time and gain of a simulation with increasing generation time.

Figure 15 depicts the average execution time of 5 runs where each simulation generates 500 elements. The process time is fixed to 60,000 ms, while the generation time between stream elements varies from 100 ms to 2,000 ms. For short generation times, almost all the processing tasks are executed when the

generation task has already finished, obtaining no gain with respect to the implementation with objects. For instance, when generating elements every 100 ms, the simulation takes 50,000 ms in total ($500 \text{ elements} \cdot 100 \text{ ms/element}$). Since the process tasks last 60,000 ms, none of them will have finished before the simulation ends; leading to almost no gain.

When increasing the generation time, more and more tasks can be executed while the generation is still active; achieving a 19% gain when generating stream elements every 500 ms. However, the gain is limited because the last generated elements are always processed when the simulation is over. Therefore, increasing the generation time from 500 ms to 2,000 ms only raises the gain from 19% to 23%.




Figure 16: Average execution time and gain of a simulation with increasing process time.

On the other hand, Figure 16 illustrates the average execution time of 5 runs that generate 500 process tasks with a fixed generation time of 100 ms and a process time varying from 5,000 ms up to 60,000 ms. Notice that the total simulation time is 50,000 ms. When the processing time is short, many tasks can be executed while the generation is still active; achieving a maximum 23% gain when the processing time is 5,000ms.

As with the previous case, when the processing time is increased, the number of tasks that can be executed while the generation is active also decreases and, thus, the gain. Also, the gain is almost zero when the processing time is big

enough (60,000ms) so that none of the process tasks will have finished before the generation ends.

6.3. Gain of removing synchronisations

Many workflows are composed of several iterative computations running simultaneously until a certain convergence criterion is met. As described in Section 5.2, this technique is useful when performing parameter sweep, cross-validation, or running the same algorithm with different initial points.




Figure 17: Parallel iterative computations. Pure task-based workflow and Hybrid Workflow shown at left and right, respectively.

To this end, each computation requires a phase at the end of each iteration to exchange information with the rest. When using pure task-based workflows, this phase requires to stop all the computations at the end of each iteration, retrieve all the states, create a task to exchange and update all the states, transfer back all the new states, and resume all the computations to the next iteration. The left task graph of Figure 17 shows an example of such workflows with two iterations of two parallel computations. The first two red tasks initialise the state of each computation, the pink tasks perform the computation of each iteration, and the blue tasks retrieve and update the state of each computation.

Conversely, when using Hybrid Workflows, each computation can exchange the information at the end of each iteration asynchronously by writing and reading the states to/from streams. This technique avoids splitting each computation into tasks, stopping and resuming each computation at every iteration, and synchronising all the computations to exchange data. The right task graph of Figure 17 depicts the equivalent Hybrid Workflow of the previous example. Each computation is run in a single task (white) that performs the state initialisation, all the iterations, and all the update phases at the end of each iteration.




Figure 18: Average execution time and gain of a simulation with an increasing number of iterations.

Using the previous examples, Figure 18 evaluates the performance gain of avoiding the synchronisation and exchange process at the end of each iteration (calculated following Equation 2). Hence, the benchmark executes the pure task-based workflow (blue) and the Hybrid Workflow (green) versions of the same workflow written in Java and using Kafka as streaming backend. Also, it is composed of two independent computations with a fixed computation per iteration (2,000 ms) and an increasing number of iterations. The results shown are the mean execution times of 5 runs of each configuration.

$$Gain = \frac{ExecutionTime_{pure\ task-based} - ExecutionTime_{hybrid}}{ExecutionTime_{pure\ task-based}} \quad (2)$$

Notice that the total gain is influenced by three factors: the removal of the synchronisation task at the end of each iteration, the cost of transferring the state between the process and the synchronisation tasks, and, the division of the state's initialisation and process. Although we have reduced the state of each computation to 24 bytes and used a single worker machine to minimise the impact of the transfer overhead, the second and third factors become important when running a small number of iterations (below 32), reaching a maximum gain of 42% when running a single iteration. For a larger number of iterations (over 32), the removal of the synchronisation becomes the main factor, and the total gain reaches a steady state with a gain around 33%.

6.4. Stream writers and readers scalability and load balance

Our prototype supports N-M streams, meaning that any stream can have an arbitrary amount of writers and readers. To evaluate the performance and load balance, we have implemented a Java application that uses a single stream and creates N writer tasks and M reader tasks. Although our writer and reader tasks use a single core, we spawn each of them in separated nodes so that the data must be transferred. In more sophisticated use cases, each task could benefit from an intra-node technology (such as OpenMP) to parallelise the processing of the stream data.




Figure 19: Average execution time and efficiency (left and right, respectively) with increasing number of readers and different number of writers.

Figure 19 depicts the average execution time (left) and the efficiency (right) of 5 runs with an increasing number of readers. Each series uses a different

number of writers, also going from 1 to 8. Also, the writers publish 100 elements in total, the size of the published objects is 24 bytes, and the time to process an element is set to 1,000ms. The efficiency is calculated using the ideal execution time as reference; i.e. the number of elements multiplied by the time to process an element and divided by the number of readers.

Since the execution time is mainly due to the processing of the elements in the reader tasks, when increasing the number of writers, there are no significant differences. However, for all the cases, increasing the number of readers significantly impacts the execution time, achieving a 4.84 speed-up with 8 readers. Furthermore, the efficiencies using 1 reader are close to the ideal (87% on average) because the only overheads are the creation of the elements, the task spawning, and the data transfers. However, when increasing the number of readers, the load imbalance significantly affects efficiency; achieving around 50% efficiency with 8 readers.




Figure 20: Number of stream elements processed per reader.

It is worth mentioning that the achieved speed-up is lower than the ideal (8) due to load imbalance. Thus, the elements processed by each reader task are not balanced since elements are assigned to the first process that requests them. Figure 20 illustrates an in-depth study of the load imbalance when running 1, 2, 4, or 8 readers. Notice that when running with 2 readers, the first reader gets almost 75% of the elements while the second one only processes 25% of the total load. The same pattern is shown when increasing the number of readers; where half of the tasks perform 70% of the total load. For instance, when running with 4 readers, 2 tasks perform 69% of the work (34.5% each), while the rest only performs 31% of the total load (15.5% each). Similarly, when running with

8 readers, 4 tasks perform 70% of the total load (17.5% each), while the other four only process 30% (7.5% each).

At its current state, the *Distributed Stream Library* does not implement any load balance technique, nor limit the number of elements retrieved by each `poll` call. As future work, since the library already stores the processes registered to each stream, it could implement some policy and send only a subset of the available elements to the requesting process rather than all of them.

6.5. Runtime overhead

To provide a more in-depth analysis of the performance of our prototype, we have compared each step of the task life-cycle when using `ObjectParameter` (OP from now on) or `StreamParameter` (SP from now on) that use `ObjectDistroStreams`. The following figures evaluate the task analysis, task scheduling, and task execution average times of 100 tasks using (a) a single object of increasing size (from 1 MB to 128 MB) or (b) an increasing number of objects (from 1 to 16 objects) of fixed size (8 MB). Both implementations are written in Java and Kafka is used as the stream backend. Regarding the task definition, notice that the OP implementation requires an `ObjectParameter` for each object sent to the task. In contrast, the SP implementation only requires a single `StreamParameter` since all the objects are sent through the stream itself.




Figure 21: Task analysis average time for one single parameter with increasing sizes (left) or increasing number of parameters (right).

Figure 21 compares the task analysis results. The task analysis is the time spent by the runtime to register the task and its parameters into the system. It is worth mentioning that increasing the object's size does not affect the anal-

ysis time in either the OP nor the SP implementations. There is, however, a difference around 0.05 ms due to the creation of internal structures to represent object parameters or stream parameters.

On the other hand, increasing the number of objects directly affects the task analysis time because the runtime needs to register each task parameter individually. For the OP implementation, each object maps to an extra task parameter, and thus, the task analysis time slightly increases when increasing the number of objects. Conversely, for the SP implementation, the stream parameter itself is not modified since we only increase the number of published objects. Hence, the task analysis time remains constant when increasing the number of objects.




Figure 22: Task scheduling average time for one single parameter with increasing sizes (left) or increasing number of parameters (right).

Figure 22 compares the task scheduling results. On the one hand, the scheduling time for both implementations varies from 2.05 ms to 2.20 ms but does not show any clear tendency to increase when increasing the object's size. On the other hand, when increasing the number of objects, the scheduling time increases for the OP implementation and remains constant for the SP implementation. This behaviour is due to the fact that the default COMPSs scheduler implements data locality and, thus, the scheduling time is proportional to the number of parameters. Similarly to the previous case, increasing the number of objects increases the number of task parameters for the OP implementation (increasing its scheduling time), but keeps a single parameter for the SP implementation (maintaining its scheduling time).




Figure 23: Task execution average time for one single parameter with increasing sizes (left) or increasing number of parameters (right).

Figure 23 compares the task execution results. The task execution time covers the transfer of all the task parameters and the task execution itself. Regarding the SP implementation, the time remains constant around 208 ms regardless of the object’s size and the number of objects because the measurement only considers the transfer of the stream object itself and the execution time of the `poll` method. It is worth mentioning that the actual transfers of the objects are done by Kafka when invoking the `publish` method on the main code, and thus, they are executed in parallel while COMPSs spawns the task in the worker machine.

Conversely, the execution time for the OP implementation increases with both the object’s size and the number of objects, since the serialisation and transfer times also increase. However, the task execution does not need to fetch the objects (the `poll` method) since all of them have already been transferred. This trade-off can be observed in the figure, where the OP implementation performs better than the SP implementation when using task parameters smaller than 48 MB and performs worse for bigger cases. Notice that only the total objects’ size is relevant since the same behaviour is shown when using a single 48 MB object or 6 objects of 8 MB each.

Since the real object transfers when using SP are executed during the `publish` method and cannot be observed measuring the task execution time, we have also measured the total execution time of the benchmark for both implementations. Figure 24 shows the total execution time with an increasing number of objects of

8 MB. In contrast to the previous plot, both implementations have an increasing execution time proportional to the objects' size. Also, the SP implementation only outperforms the OP implementation when using more than 12 objects.

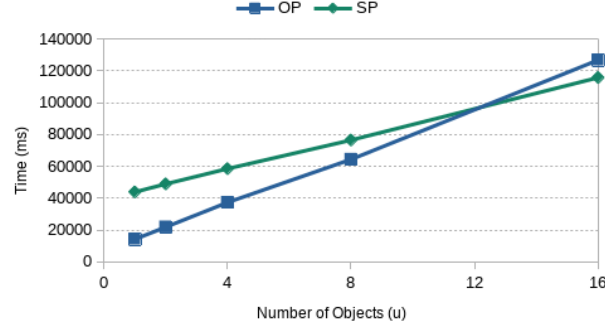


Figure 24: Total execution time with increasing number of parameters.

To conclude, since there are no major differences regarding the task analysis time nor the task scheduling time, we can safely assume that the use of streams instead of regular objects is recommended when the total size of the task parameters exceeds 48 MB and there are more than 12 objects published to the stream.

7. Conclusion and Future Work

This paper demonstrates that task-based workflows and dataflows can be integrated into a single programming model to better-cover the needs of the new Data Science workflows. Using Hybrid Workflows, developers can build complex pipelines with different approaches at many levels using a single framework.

The proposed solution relies on the *DistroStream* concept: a generic API used by applications to handle stream accesses homogeneously regardless of the software backing it. Two implementations provide the specific logic to support object and file streams. The first one, *ObjectDistroStream*, is built on top of Kafka to enable object streams. The second one, *FileDistroStream*, monitors the creation of files inside a directory, sends the file locations through the stream, and relies on a distributed file system to share the file content.

The *DistroStream* API and both implementations are part of the *DistroStream-Lib*, which also provides the *DistroStream Client* and the *DistroStream Server*. While the client acts as a broker on behalf of the application and interacts with the corresponding backend, the server manages the streams’ metadata and coordinates the accesses.

By integrating the *DistroStreamLib* into a task-based Workflow Manager, its programming model can easily support Hybrid Workflows. The described prototype extends COMPSs to enable tasks with continuous input and output data by providing a new annotation for stream parameters. Implementing the handling of such stream-type values lead to some modifications on the Task Analyser and Task Scheduler components of the runtime. Using the *DistroStreamLib* also implied changes at deployment time since its components need to be spawned along with COMPSs. On the one hand, the COMPSs master hosts the *DistroStream Server*, the required stream backend, and a *DistroStream Client* to handle stream accesses on the application’s main code. On the other hand, each COMPSs worker contains a *DistroStream Client* that performs the stream accesses on tasks. Although the described prototype only builds on COMPSs, it can be used as an implementation reference for any other existing task-based framework.

This paper also presents four use cases illustrating the new capabilities that the users may identify in their workflows to benefit from the use of Hybrid Workflows. On the one hand, streams can be internal or external to the application and can be used to communicate continuous data or control data. On the other hand, streams can be accessed inside the main code, native tasks (i.e., Java or Python), or non-native tasks (i.e., MPI, binaries, and nested COMPSs workflows). Furthermore, the *Distributed Stream Library* supports the one to many, many to one, and many to many scenarios transparently, and allows to configure the consumer mode to process the data at least once, at most once, or exactly once when using many consumers.

The evaluation demonstrates the benefit of processing data continuously as it is generated; achieving a 23% gain with the right generation and process times

and resources. Also, using streams as control mechanism enabled the removal of synchronisation points when running several parallel algorithms, leading to a 33% gain when running more than 32 iterations. Finally, an in-depth analysis of the runtime’s performance shows that there are no major differences regarding the task analysis time nor the task scheduling time when using streams or object tasks, and that the use of streams is recommended when the total size of the task parameters exceeds 48 MB using more than 12 objects.

Although the solution is fully functional, some improvements can be made. Regarding the *DistroStream* implementations, we plan to extend the *FileDistroStream* to support shared disks with different mount-points. On the other hand, we will add new *ObjectDistroStream*’s backend implementations (apart from Kafka), so that users can choose between them without changing the application’s code. Envisaging that a single *DistroStream Server* could become a bottleneck when managing several applications involving a large number of cores, we consider replacing the client-server architecture by a peer-to-peer approach. Finally, by highlighting the benefits from hybrid flows, we expect to attract real-world applications to elevate our evaluation to more complex use cases.

Acknowledgements

This work has been supported by the Spanish Government (contracts SEV2015-0493 and TIN2015-65316-P), by Generalitat de Catalunya (contract 2014-SGR-1051), and by the European Commission through the Horizon 2020 Research and Innovation program under contract 730929 (MF2C project). Cristian Ramon-Cortes predoctoral contract is financed by the Spanish Government under the contract BES-2016-076791.

References

References

- [1] P. A. Damkroger, The Intersection of AI, HPC and HPDA: How Next-Generation Workflows Will Drive Tomorrows Breakthroughs, Accessed:

August 07, 2019 (2018).

URL <https://www.top500.org/>

- [2] J. Kreps, N. Narkhede, J. Rao, et al., Kafka: A distributed messaging system for log processing, in: Proceedings of the NetDB, 2011, pp. 1–7.
URL <http://people.csail.mit.edu/matei/courses/2015/6.S897/readings/kafka.pdf>
- [3] R. M. Badia, et al., COMP superscalar, an interoperable programming framework, *SoftwareX* 3 (2015) 32–36.
URL <https://doi.org/10.1016/j.softx.2015.10.004>
- [4] Barcelona Supercomputing Center (BSC), COMP Superscalar — BSC-CNS, Accessed: March 26, 2019.
URL <http://compss.bsc.es>
- [5] A. Jain, S. P. Ong, W. Chen, B. Medasani, X. Qu, M. Kocher, M. Brafman, G. Petretto, G.-M. Rignanesi, G. Hautier, et al., Fireworks: A dynamic workflow system designed for high-throughput applications, *Concurrency and Computation: Practice and Experience* 27 (17) (2015) 5037–5059.
- [6] A. Jain, Introduction to FireWorks (workflow software) - FireWorks 1.8.7 documentation, Accessed: March 26, 2019.
URL <https://materialsproject.github.io/fireworks/#>
- [7] K. Wolstencroft, R. Haines, D. Fellows, A. Williams, D. Withers, S. Owen, S. Soiland-Reyes, I. Dunlop, A. Nenadic, P. Fisher, J. Bhagat, K. Belhajjame, F. Bacall, A. Hardisty, A. Nieva de la Hidalga, M. P. Balcazar Vargas, S. Sufi, C. Goble, The Taverna workflow suite: designing and executing workflows of Web Services on the desktop, web or in the cloud, *Nucleic Acids Research* 41 (W1) (2013) W557–W561. doi:10.1093/nar/gkt328.
URL <https://doi.org/10.1093/nar/gkt328>
- [8] Taverna Committers, Apache Taverna, Accessed: April 1, 2019.
URL <https://taverna.incubator.apache.org/>

- [9] B. Ludäscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger, M. Jones, E. A. Lee, J. Tao, Y. Zhao, Scientific workflow management and the kepler system, *Concurrency and Computation: Practice and Experience* 18 (10) (2006) 1039–1065.
- [10] UC Davis and UC Santa Barbara and UC San Diego, The Kepler Project, Accessed: April 1, 2019.
URL <https://kepler-project.org/>
- [11] E. Afgan, D. Baker, M. van den Beek, D. Blankenberg, D. Bouvier, M. Čech, J. Chilton, D. Clements, N. Coraor, C. Eberhard, B. Grüning, A. Guerler, J. Hillman-Jackson, G. Von Kuster, E. Rasche, N. Soranzo, N. Turaga, J. Taylor, A. Nekrutenko, J. Goecks, The Galaxy platform for accessible, reproducible and collaborative biomedical analyses: 2016 update, *Nucleic Acids Res.* 44 (W1) (2016) W3–W10.
- [12] Galaxy Team, Galaxy, Accessed: March 26, 2019.
URL <https://usegalaxy.org/>
- [13] M. Abouelhoda, S. A. Issa, M. Ghanem, Tavaxy: Integrating Taverna and Galaxy workflows with cloud computing support, *BMC bioinformatics* 13 (1) (2012) 77.
- [14] P. Di Tommaso, M. Chatzou, E. W. Floden, P. P. Barja, E. Palumbo, C. Notredame, Nextflow enables reproducible computational workflows, *Nature biotechnology* 35 (4) (2017) 316.
- [15] Barcelona Centre for Genomic Regulation (CRG), Nextflow: A DSL for parallel and scalable computational pipelines, Accessed: March 26, 2019.
URL <https://www.nextflow.io/>
- [16] M. Wilde, M. Hategan, J. M. Wozniak, B. Clifford, D. S. Katz, I. Foster, Swift: A language for distributed parallel scripting, *Parallel Computing* 37 (9) (2011) 633 – 652. doi:<https://doi.org/10.1016/j.parco.2011.05.005>.

URL <http://www.sciencedirect.com/science/article/pii/S0167819111000524>

- [17] Swift Project Team, The Swift Parallel Scripting Language, Accessed: March 26, 2019.
URL <http://swift-lang.org/main/>
- [18] University of Chicago, Parsl: Parallel Scripting in Python, Accessed: March 26, 2019.
URL <http://parsl-project.org/>
- [19] Dask Development Team, Dask: Library for dynamic task scheduling (2016).
URL <https://dask.org>
- [20] F. Lordan, R. M. Badia, et al., ServiceSs: an interoperable programming framework for the Cloud, *Journal of Grid Computing* 12 (1) (2014) 67–91.
URL <https://digital.csic.es/handle/10261/132141>
- [21] Apache Flink Contributors, Apache Flink, Accessed: August 9, 2019 (2014).
URL <https://flink.apache.org/>
- [22] Apache Samza Contributors, Apache Samza, Accessed: August 9, 2019 (2019).
URL <http://samza.apache.org/>
- [23] A. Toshniwal, et al., Storm@ twitter, in: *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, ACM, 2014, pp. 147–156.
- [24] S. Kulkarni, et al., Twitter heron: Stream processing at scale, in: *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, ACM, 2015, pp. 239–250.

- [25] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, I. Stoica, Spark: Cluster computing with working sets., HotCloud 10 (10-10) (2010) 95.
- [26] J. Dean, S. Ghemawat, Mapreduce: Simplified data processing on large clusters, Proceedings of the 6th Conference on Symposium on Operating Systems Design and Implementation 6 (2004) 10–10.
URL <http://dl.acm.org/citation.cfm?id=1251254.1251264>
- [27] M. Zaharia, et al., Discretized Streams: An Efficient and Fault-Tolerant Model for Stream Processing on Large Clusters, HotCloud 12 (2012) 10–16.
- [28] R. Filgueira, R. F. d. Silva, A. Krause, E. Deelman, M. Atkinson, Asterism: Pegasus and Dispel4py Hybrid Workflows for Data-Intensive Science, in: 2016 Seventh International Workshop on Data-Intensive Computing in the Clouds (DataCloud), 2016, pp. 1–8. doi:10.1109/DataCloud.2016.004.
- [29] E. Tejedor, R. M. Badia, J. Labarta, et al., PyCOMPSs: Parallel computational workflows in Python, The International Journal of High Performance Computing Applications (IJHPCA) 31 (2017) 66–82.
URL <http://dx.doi.org/10.1177/1094342015594678>
- [30] J. Ejarque, M. Domnguez, R. M. Badia, A hierarchic task-based programming model for distributed heterogeneous computing, The International Journal of High Performance Computing Applications 0 (0) (0) 1094342019845438. arXiv:<https://doi.org/10.1177/1094342019845438>, doi:10.1177/1094342019845438.
URL <https://doi.org/10.1177/1094342019845438>
- [31] Barcelona Supercomputing Center (BSC), Extrae Tool, Accessed: July 31, 2019.
URL <https://tools.bsc.es/extrae>
- [32] Barcelona Supercomputing Center (BSC), Paraver Tool, Accessed: July

31, 2019.

URL <https://tools.bsc.es/paraver>

- [33] S. Liang, Java Native Interface: Programmer's Guide and Reference, 1st Edition, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.
- [34] P. Hunt, M. Konar, F. P. Junqueira, B. Reed, ZooKeeper: Wait-free Coordination for Internet-scale Systems, in: USENIX annual technical conference, Vol. 8, Boston, MA, USA, 2010, p. 14.
- [35] Apache ZooKeeper Contributors, Apache ZooKeeper, Accessed: July 31, 2019.
URL <https://zookeeper.apache.org/>
- [36] Oracle and/or its affiliates, Java NIO, Accessed: January 20, 2020 (2010).
URL <https://docs.oracle.com/javase/1.5.0/docs/guide/nio/index.html>
- [37] Barcelona Supercomputing Center (BSC), MareNostrum IV Technical Information, Accessed: July 31, 2019.
URL <https://www.bsc.es/marenostrum/marenostrum/technical-information>
- [38] Workflows and Distributed Computing - Barcelona Supercomputing Center (BSC), DistroStream Library GitHub, Accessed: September 3, 2019.
URL <https://github.com/bsc-wdc/distro-stream-lib>
- [39] Workflows and Distributed Computing - Barcelona Supercomputing Center (BSC), COMPSs GitHub, Accessed: July 31, 2019.
URL <https://github.com/bsc-wdc/compss>
- [40] Apache Kafka Contributors, Apache Kafka, Accessed: July 12, 2019.
URL <https://github.com/apache/kafka>