



COMP SUPERSCALAR

---

# User Manual

Application development guide

---

VERSION: 2.5

June 14, 2019



**Barcelona  
Supercomputing  
Center**

*Centro Nacional de Supercomputación*

This manual only provides information about the development of COMPSs applications. Specifically, it details the programming model features available in Java, Python and C/C++ languages.

For an extensive list of COMPSs application examples (codes, execution commands, results, logs, etc.) please refer to the *COMPSs Sample Applications* guide at <http://compss.bsc.es/> .

For information about the installation process please refer to the *COMPSs Installation Guide* available at <http://compss.bsc.es/> .

For further information about the application execution please refer to the *COMPSs User Manual: Application execution guide* available at <http://compss.bsc.es/> .

# Contents

<b>1</b>	<b>COMP Superscalar (COMPSs)</b>	<b>1</b>
<b>2</b>	<b>Java</b>	<b>2</b>
2.1	Programming Model . . . . .	2
2.1.1	Main application code . . . . .	2
2.1.2	Remote methods code . . . . .	3
2.1.3	Java annotated interface . . . . .	4
2.1.4	Alternative method implementations . . . . .	7
2.1.5	Java API calls . . . . .	8
2.2	Application Compilation . . . . .	9
2.3	Application Execution . . . . .	10
<b>3</b>	<b>Python Binding</b>	<b>11</b>
3.1	Programming Model . . . . .	11
3.1.1	Task Selection . . . . .	11
3.1.1.1	Scheduler hints . . . . .	15
3.1.1.2	Other task types . . . . .	16
3.1.1.2.1	Binary decorator . . . . .	17
3.1.1.2.2	OmpSs decorator . . . . .	18
3.1.1.2.3	MPI decorator . . . . .	19
3.1.1.2.4	COMPSs decorator . . . . .	19
3.1.1.2.5	Multinode decorator . . . . .	19
3.1.1.2.6	Parameters summary . . . . .	20
3.1.2	Constraints . . . . .	22
3.1.3	Implements . . . . .	23
3.1.4	Main Program . . . . .	24
3.1.5	Important Notes . . . . .	26
3.2	Application Execution . . . . .	28
3.2.1	Environment . . . . .	28
3.2.2	Command . . . . .	28
3.3	Development with Jupyter notebook . . . . .	28
3.3.1	Environment . . . . .	28
3.3.2	API calls . . . . .	29
3.3.3	Application execution . . . . .	32
3.4	Integration with Numba . . . . .	32
3.4.1	Basic usage . . . . .	32
3.4.2	Advanced usage . . . . .	32
<b>4</b>	<b>C/C++ Binding</b>	<b>34</b>
4.1	Programming Model . . . . .	34
4.1.1	Task Selection . . . . .	34
4.1.2	Main Program . . . . .	34
4.1.3	Binding API . . . . .	36
4.1.4	Functions file . . . . .	36

4.1.5	Additional source Files . . . . .	37
4.1.6	Class Serialization . . . . .	37
4.1.7	Method - Task . . . . .	38
4.1.8	Task Constraints . . . . .	39
4.1.9	Task Versions . . . . .	39
4.2	Application Compilation . . . . .	39
4.2.1	Single architecture . . . . .	40
4.2.2	Multiple architectures . . . . .	41
4.3	Application Execution . . . . .	43
4.4	Task Dependency Graph . . . . .	43
<b>5</b>	<b>Constraints</b>	<b>45</b>
<b>6</b>	<b>Known Limitations</b>	<b>48</b>

## List of Figures

1	Matmul Execution Graph. . . . .	44
---	---------------------------------	----

## List of Tables

1	Arguments of the <i>@task</i> decorator. . . . .	16
2	<i>@binary</i> , <i>@ompss</i> , <i>@mpi</i> , <i>@compss</i> and <i>@multinode</i> decorators supported parameters. . . . .	20
3	Supported Streams for the <i>@binary</i> , <i>@ompss</i> and <i>@mpi</i> decorators . . . . .	21
4	File parameters definition shortcuts . . . . .	22
5	COMPSs Python API functions. . . . .	25
6	PyCOMPSs start function for Jupyter notebook . . . . .	31
7	PyCOMPSs stop function for Jupyter notebook . . . . .	31
8	Arguments of the <i>@constraint</i> decorator . . . . .	46
9	Arguments of the <i>@Processor</i> decorator . . . . .	47

# 1 COMP Superscalar (COMPSs)

COMP Superscalar (COMPSs) is a programming model which aims to ease the development of applications for distributed infrastructures, such as Clusters, Grids and Clouds. COMP Superscalar also features a runtime system that exploits the inherent parallelism of applications at execution time.

For the sake of programming productivity, the COMPSs model has four key characteristics:

- **Sequential programming:** COMPSs programmers do not need to deal with the typical duties of parallelization and distribution, such as thread creation and synchronization, data distribution, messaging or fault tolerance. thus eliminating most of the difficulties of concurrent/distributed programming. A task is a method or a service called from the application code that is intended to be spawned asynchronously and possibly run in parallel with other tasks on a set of resources, instead of locally and sequentially.
- **Infrastructure unaware:** COMPSs offers a model that abstracts the application from the underlying infrastructure. Hence, COMPSs programs do not include any detail that could tie them to a particular platform, like deployment or resource management. This makes applications portable between infrastructures with diverse characteristics.
- **Standard programming languages:** COMPSs natively supports Java applications, but also offers language bindings for Python and C/C++ applications.
- **No APIs:** In the case of COMPSs applications in Java, the model does not require to use any special API call, pragma or construct in the application; everything is standard Java syntax and libraries. As regards the Python and C/C++ bindings, a small set of API calls should be used on the COMPSs applications.

## 2 Java

This section illustrates the steps to develop a Java COMPSs application, to compile and to execute it. The *Simple* application will be used as reference code. The user is required to select a set of methods, invoked in the sequential application, that will be run as remote tasks on the available resources.

### 2.1 Programming Model

A COMPSs application is composed of three parts:

- **Main application code:** the code that is executed sequentially and contains the calls to the user-selected methods that will be executed by the COMPSs runtime as asynchronous parallel tasks.
- **Remote methods code:** the implementation of the tasks.
- **Java annotated interface:** It declares the methods to be run as remote tasks along with metadata information needed by the runtime to properly schedule the tasks.

The main application file name has to be the same of the main class and starts with capital letter, in this case it is **Simple.java**. The Java annotated interface filename is *application name+Itf.java*, in this case it is **SimpleItf.java**. And the code that implements the remote tasks is defined in the *application name + Impl.java* file, in this case it is **SimpleImpl.java**.

All code examples are in the `/home/compss/tutorial_apps/java/` folder of the development environment.

#### 2.1.1 Main application code

In COMPSs the user's application code is kept unchanged, no API calls need to be included in the main application code in order to run the selected tasks on the nodes.

The COMPSs runtime is in charge of replacing the invocations to the user-selected methods with the creation of remote tasks also taking care of the access to files where required. Let's consider the Simple application example that takes an integer as input parameter and increases it by one unit.

The main application code of Simple app (**Simple.java**) is executed sequentially until the call to the **increment()** method. COMPSs, as mentioned above, replaces the call to this method with the generation of a remote task that will be executed on an available node.

```
package simple;

import java.io.FileInputStream;
import java.io.FileOutputStream;
import java.io.IOException;
import simple.SimpleImpl;

public class Simple {
```



```

public static void main(String[] args) {
    String counterName = "counter";
    int initialValue = args[0];

    //-----//
    //Creation of the file which will contain the counter variable//
    //-----//
    try {
        FileOutputStream fos = new FileOutputStream(counterName);
        fos.write(initialValue);
        System.out.println("Initial counter value is " + initialValue);
        fos.close();
    }catch(IOException ioe) {
        ioe.printStackTrace();
    }

    //-----//
    //           Execution of the program           //
    //-----//
    SimpleImpl.increment(counterName);

    //-----//
    //   Reading from an object stored in a File   //
    //-----//
    try {
        FileInputStream fis = new FileInputStream(counterName);
        System.out.println("Final counter value is " + fis.read());
        fis.close();
    }catch(IOException ioe) {
        ioe.printStackTrace();
    }
}
}

```

### 2.1.2 Remote methods code

The following code contains the implementation of the remote method of the *Simple* application (**SimpleImpl.java**) that will be executed remotely by COMPSs.

```

package simple;

import java.io.FileInputStream;
import java.io.FileOutputStream;
import java.io.IOException;
import java.io.FileNotFoundException;

public class SimpleImpl {
    public static void increment(String counterFile) {
        try{
            FileInputStream fis = new FileInputStream(counterFile);
            int count = fis.read();
            fis.close();
            FileOutputStream fos = new FileOutputStream(counterFile);
            fos.write(++count);
            fos.close();
        }catch(FileNotFoundException fnfe){
            fnfe.printStackTrace();
        }catch(IOException ioe){
            ioe.printStackTrace();
        }
    }
}
}

```

### 2.1.3 Java annotated interface

The Java interface is used to declare the methods to be executed remotely along with Java annotations that specify the necessary metadata about the tasks. The metadata can be of three different types:

1. For each parameter of a method, the data type (currently *File* type, primitive types and the *String* type are supported) and its directions (IN, OUT, INOUT or CONCURRENT).
2. The Java class that contains the code of the method.
3. The constraints that a given resource must fulfill to execute the method, such as the number of processors or main memory size.

A complete and detailed explanation of the usage of the metadata includes:

- **Method-level Metadata:** for each selected method, the following metadata has to be defined:
  - **@Method:** Defines the Java method as a task
    - \* **declaringClass** (Mandatory) String specifying the class that implements the Java method.
    - \* **targetDirection** This field specifies the direction of the target object of an object method. It can be defined as: "INOUT" (default value) if the method modifies the target object, "CONCURRENT" if this object modification can be done concurrently, or "IN" if the method does not modify the target object.
    - \* **priority** "true" if the task takes priority and "false" otherwise. This parameter is used by the COMPSs scheduler (it is a String not a Java boolean).
    - \* **onFailure** Expected behaviour if the task fails. *OnFailure.RETRY* (default value) makes the task be executed again, *OnFailure.CANCEL\_SUCCESORS* ignores the failure and cancels the sucesor tasks, *OnFailure.FAIL* stops the whole application in a save mode once a task fails or *OnFailure.IGNORE* ignores the failure and continues with normal runtime execution.
  - **@Binary:** Defines the Java method as a binary invokation
    - \* **binary** (Mandatory) String defining the full path of the binary that must be executed.
    - \* **workingDir** Full path of the binary working directory inside the COMPSs Worker.
    - \* **priority** "true" if the task takes priority and "false" otherwise. This parameter is used by the COMPSs scheduler (it is a String not a Java boolean).
  - **@MPI:** Defines the Java method as a MPI invokation
    - \* **mpiRunner** (Mandatory) String defining the mpi runner command.

- \* **binary** (Mandatory) String defining the full path of the binary that must be executed.
  - \* **computingNodes** String defining the number of computing nodes reserved for the MPI execution (only a single node is reserved by default).
  - \* **workingDir** Full path of the binary working directory inside the COMPSs Worker.
  - \* **priority** "true" if the task takes priority and "false" otherwise. This parameter is used by the COMPSs scheduler (it is a String not a Java boolean).
- **@OmpSs:** Defines the Java method as a OmpSs invocation
    - \* **binary** (Mandatory) String defining the full path of the binary that must be executed.
    - \* **workingDir** Full path of the binary working directory inside the COMPSs Worker.
    - \* **priority** "true" if the task takes priority and "false" otherwise. This parameter is used by the COMPSs scheduler (it is a String not a Java boolean).
  - **@Constraints:** The user can specify the capabilities that a resource must have in order to run a method. For example, in a cloud execution the COMPSs runtime creates a VM that fulfils the specified requirements in order to perform the execution. A full description of the supported constraints can be found in Table 8 in Section 5.
  - **@SchedulerHints:** It specifies the class that implements the method.
    - \* **isReplicated** "true" if the method must be executed in all the worker nodes when invoked from the main application (it is a String not a Java boolean).
    - \* **isDistributed** "true" if the method must be scheduled in a forced round robin among the available resources (it is a String not a Java boolean).
- **Parameter-level Metadata (@Parameter):** for each parameter and method, the user must define:
    - **Direction:** *Direction.IN*, *Direction.INOUT*, *Direction.OUT* or *Direction.CONCURRENT*
    - **Type:** COMPSs supports the following types for task parameters:
      - \* **Basic types:** *Type.BOOLEAN*, *Type.CHAR*, *Type.BYTE*, *Type.SHORT*, *Type.INT*, *Type.LONG*, *Type.FLOAT*, *Type.DOUBLE*. They can only have **IN** direction, since primitive types in Java are always passed by value.
      - \* **String:** *Type.STRING*. It can only have **IN** direction, since Java Strings are immutable.
      - \* **File:** *Type.FILE*. It can have any direction (**IN**, **OUT**, **INOUT** or **CONCURRENT**). The real Java type associated with a **FILE** parameter is a String that contains the path to the file. However, if the user specifies a parameter as a **FILE**, COMPSs will treat it as such.

- \* **Object:** *Type.Object*. It can have any direction (IN, OUT, INOUT or CONCURRENT).
- **Return type:** Any object or a generic class object. In this case the direction is always OUT.

Basic types are also supported as return types. However, we do not recommend to use them because they cause an implicit synchronization

- **Stream:** For non-native tasks (binaries, MPI, and OmpSs) COMPSs supports the automatic redirection of the Linux streams by specifying `Stream.STDIN`, `Stream.STDOUT` or `Stream.STDERR`. Notice that any parameter annotated with the stream annotation must be of type *Type.FILE*, and with direction *Direction.IN* for *Stream.STDIN* or *Direction.OUT/ Direction.INOUT* for *Stream.STDOUT* and *Stream.STDERR*.
- **Prefix:** For non-native tasks (binaries, MPI, and OmpSs) COMPSs allows to prepend a constant String to the parameter value to use the Linux joint-prefixes as parameters of the binary execution.
- **Service-level Metadata:** for each selected service, the following metadata has to be defined:
  - **@Service:** Mandatory. It specifies the service properties.
    - \* **namespace** Mandatory. Service namespace
    - \* **name** Mandatory. Service name.
    - \* **port** Mandatory. Service port.
    - \* **operation** Operation type.
    - \* **priority** "true" if the service takes priority, "false" otherwise. This parameter is used by the COMPSs scheduler (it is a String not a Java boolean).

The Java annotated interface of the Simple app example (`SimpleItf.java`) includes the description of the *Increment()* method metadata. The method interface contains a single input parameter, a string containing a path to the file counterFile. In this example there are constraints on the minimum number of processors and minimum memory size needed to run the method.

```
package simple;

import es.bsc.compss.types.annotations.Constraints;
import es.bsc.compss.types.annotations.task.Method;
import es.bsc.compss.types.annotations.Parameter;
import es.bsc.compss.types.annotations.parameter.Direction;
import es.bsc.compss.types.annotations.parameter.Type;

public interface SimpleItf {
```

```

@Constraints(computingUnits = "1", memorySize = "0.3")
@Method(declaringClass = "simple.SimpleImpl")
void increment(
    @Parameter(type = Type.FILE, direction = Direction.INOUT)
    String file
);
}

```

#### 2.1.4 Alternative method implementations

Since version 1.2, the COMPSs programming model allows developers to define sets of alternative implementations of the same method in the Java annotated interface. The following code depicts an example where the developer sorts an integer array using two different methods: merge sort and quick sort that are respectively hosted in the *packagepath.Mergesort* and *packagepath.Quicksort* classes.

```

@Method(declaringClass = "packagepath.Mergesort")
@Method(declaringClass = "packagepath.Quicksort")
void sort(
    @Parameter(type = Type.OBJECT, direction = Direction.INOUT)
    int[] array
);

```

As depicted in the example, the name and parameters of all the implementations must coincide; the only difference is the class where the method is implemented. This is reflected in the attribute *declaringClass* of the *@Method* annotation. Instead of stating that the method is implemented in a single class, the programmer can define several instances of the *@Method* annotation with different declaring classes.

As independent remote methods, the sets of equivalent methods might have common restrictions to be fulfilled by the resource hosting the execution. Or even, each implementation can have specific constraints. Through the *@Constraints* annotation, developers can specify the common constraints for a whole set of methods. In the following example only one core is required to run the method of both sorting algorithms.

```

@Constraints(computingUnits = "1")
@Method(declaringClass = "packagepath.Mergesort")
@Method(declaringClass = "packagepath.Quicksort")
void sort(
    @Parameter(type = Type.OBJECT, direction = Direction.INOUT)
    int[] array
);

```

However, these sorting algorithms have different memory consumption, thus each algorithm might require a specific amount of memory and that should be stated in the implementation constraints. For this purpose, the developer can add a *@Constraints* annotation inside each *@Method* annotation containing the specific constraints for that implementation. Since the Mergesort has a higher memory consumption than the quicksort, the following example sets a requirement of 1 core and 2GB of memory for the mergesort implementation and 1 core and 500MB of memory for the quicksort.

```

@Constraints(computingUnits = "1")
@Method(declaringClass = "packagepath.Mergesort", constraints = @Constraints(memorySize = "2.0"))
@Method(declaringClass = "packagepath.Quicksort", constraints = @Constraints(memorySize = "0.5"))
void sort(
    @Parameter(type = Type.OBJECT, direction = Direction.INOUT)
    int[] array
);

```

### 2.1.5 Java API calls

COMPSs also provides a explicit synchronization call, namely *barrier*, which can be used through the COMPSs Java API. The use of *barrier* forces to wait for all tasks that have been submitted before the barrier is called. When all tasks submitted before the *barrier* have finished, the execution continues.

```

import es.bsc.compss.api.COMPSs;

public class Main {
    public static void main(String[] args) {
        // Setup counterName1 and counterName2 files
        // Execute task increment 1
        SimpleImpl.increment(counterName1);
        // API Call to wait for all tasks
        COMPSs.barrier();
        // Execute task increment 2
        SimpleImpl.increment(counterName2);
    }
}

```

When an object is used in a task, COMPSs runtime store the references of these object in the runtime data structures and generate replicas and versions in remote workers. COMPSs is automatically removing these replicas for obsolete versions. However, the reference of the last version of these objects could be stored in the runtime data-structures preventing the garbage collector to remove it when there are no references in the main code. To avoid this situation, developers can indicate the runtime that an object is not going to use any more by calling the *deregisterObject* API call. The following code snippet shows a usage example of this API call.

```

import es.bsc.compss.api.COMPSs;

public class Main {
    public static void main(String[] args) {
        final int ITERATIONS = 10;
        for (int i = 0; i < ITERATIONS; ++i) {
            Dummy d = new Dummy(d);
            TaskImpl.task(d);
            /*Allows garbage collector to delete the
            object from memory when the task is finished */
            COMPSs.deregisterObject((Object) d);
        }
    }
}

```

To synchronize files, the *getFile* API call synchronizes a file, returning the last version of file with its original name. The code below contains an example of its usage.

```
import es.bsc.compss.api.COMPSs;

public class Main {
    public static void main(String[] args) {
        for (int i=0; i<1; i++) {
            TaskImpl.task(FILE_NAME, i);
        }
        /*Waits until all tasks have finished and
        synchronizes the file with its last version*/
        COMPSs.getFile(FILE_NAME);
    }
}
```

## 2.2 Application Compilation

A COMPSs Java application needs to be packaged in a *jar* file containing the class files of the main code, of the methods implementations and of the *Itf* annotation. Next we provide a set of commands to compile the Java Simple application detailed at the *COMPSs Sample Applications* available at our website <http://compss.bsc.es> .

```
compss@bsc:~$ cd tutorial_apps/java/simple/src/main/java/simple/
compss@bsc:~/tutorial_apps/java/simple/src/main/java/simple$ javac *.java
compss@bsc:~/tutorial_apps/java/simple/src/main/java/simple$ cd ..
compss@bsc:~/tutorial_apps/java/simple/src/main/java$ jar cf simple.jar simple/
compss@bsc:~/tutorial_apps/java/simple/src/main/java$ mv ./simple.jar ../../../jar/
```

In order to properly compile the code, the CLASSPATH variable has to contain the path of the *compss-engine.jar* package. The default COMPSs installation automatically add this package to the CLASSPATH; please check that your environment variable CLASSPATH contains the *compss-engine.jar* location by running the following command:

```
$ echo $CLASSPATH | grep compss-engine
```

If the result of the previous command is empty it means that you are missing the *compss-engine.jar* package in your classpath. We recommend to automatically load the variable by editing the *.bashrc* file:

```
$ echo "# COMPSs variables for Java compilation" >> ~/.bashrc
$ echo "export CLASSPATH=$CLASSPATH:/opt/COMPSs/Runtime/compss-engine.jar" >> ~/.bashrc
```

If you are using an IDE (such as Eclipse or NetBeans) we recommend you to add the *compss-engine.jar* file as an external file to the project. The *compss-engine.jar* file is available at your current COMPSs installation under the following path:

```
/opt/COMPSs/Runtime/compss-engine.jar
```

Please notice that if you have performed a custom installation, the location of the package can be different.

An Integrated Development Environment for Eclipse is also available to simplify the development, compilation, deployment and execution COMPSs applications. For further information about the *COMPSs IDE* please refer to the *COMPSs IDE User Guide* available at <http://compss.bsc.es> .

## 2.3 Application Execution

A Java COMPSs application is executed through the *runcompss* script. An example of an invocation of the script is:

```
compss@bsc:~$ runcompss --classpath=/home/compss/tutorial_apps/java/simple/jar/simple.jar
simple.Simple 1
```

A comprehensive description of the *runcompss* command is available in the *COMPSs User Manual: Application Execution* document available at <http://compss.bsc.es> .

In addition to Java, COMPSs supports the execution of applications written in other languages by means of bindings. A binding manages the interaction of the no-Java application with the COMPSs Java runtime, providing the necessary language translation.

The next sections describe the Python and C/C++ language bindings offered by COMPSs.



## 3 Python Binding

COMPSs features a binding for Python 2 and 3 applications. The next subsections explain how to program a Python application for COMPSs and a brief overview on how to execute it.

### 3.1 Programming Model

#### 3.1.1 Task Selection

As in the case of Java, a COMPSs Python application is a Python sequential program that contains calls to tasks. In particular, the user can select as a task:

- Functions
- Instance methods: methods invoked on objects.
- Class methods: static methods belonging to a class.

The task definition in Python is done by means of Python decorators instead of an annotated interface. In particular, the user needs to add a `@task` decorator that describes the task before the definition of the function/method.

As an example, let us assume that the application calls a function `func`, which receives a file path (string parameter) and an integer parameter. The code of `func` updates the file.

```
def func(file_path, value):
    # update the file 'file_path'

if __name__=='__main__':
    my_file = '/tmp/sample_file.txt'
    func(my_file, 1)
```

Please, note that the main code is defined within `if __name__=='__main__':`. A better alternative would be to define the main code within a function and invoke it from the `if __name__=='__main__':`.

In order to select `func` as a task, the corresponding `@task` decorator needs to be placed right before the definition of the function, providing some metadata about the parameters of that function. The `@task` decorator has to be imported from the `pycomps` library:

```
from pycomps.api.task import task
```

The metadata corresponding to a parameter is specified as an argument of the decorator, whose name is the formal parameter's name and whose value defines the type and direction of the parameter. The parameter types and directions can be:

- Types: *primitive types* (integer, long, float, boolean), *strings*, *objects* (instances of user-defined classes, dictionaries, lists, tuples, complex numbers) and *files* are supported.

- Direction: it can be read-only (*IN* - default), read-write (*INOUT*), write-only (*OUT*) or in some cases concurrent (*CONCURRENT*).

COMPSs is able to automatically infer the parameter type for primitive types, strings and objects, while the user needs to specify it for files. On the other hand, the direction is only mandatory for *INOUT* and *OUT* parameters. Thus, when defining the parameter metadata in the `@task` decorator, the user has the following options:

- *IN*: the parameter is read-only. The type will be inferred.
- *INOUT*: the parameter is read-write. The type will be inferred.
- *OUT*: the parameter is write-only. The type will be inferred.
- *CONCURRENT*: the parameter is read-write with concurrent acces. The type will be inferred.
- *FILE/FILE\_IN*: the parameter is a file. The direction is assumed to be *IN*.
- *FILE\_INOUT*: the parameter is a read-write file.
- *FILE\_OUT*: the parameter is a write-only file.
- *FILE\_CONCURRENT*: the parameter is a concurrent read-write file.
- *COLLECTION\_IN*: the parameter is read-only collection.
- *COLLECTION\_INOUT*: the parameter is read-write collection.

Consequently, please note that in the following cases there is no need to include an argument in the `@task` decorator for a given task parameter:

- Parameters of primitive types (integer, long, float, boolean) and strings: the type of these parameters can be automatically inferred by COMPSs, and their direction is always *IN*.
- Read-only object parameters: the type of the parameter is automatically inferred, and the direction defaults to *IN*.

The parameter metadata is available from the `pycomps` library:

```
from pycomps.api.parameter import *
```

Continuing with the example, in the following code snippet the decorator specifies that `func` has a parameter called `f`, of type *FILE* and *INOUT* direction. Note how the second parameter, `i`, does not need to be specified, since its type (integer) and direction (*IN*) are automatically inferred by COMPSs.

```

from pycompss.api.task import task      # Import @task decorator
from pycompss.api.parameter import *    # Import parameter metadata for the @task decorator

@task (f=FILE_INOUT)
def func(f, i):
    fd = open(f, 'r+')
    ...

```

The user can also define that the access to a parameter is concurrent with *CONCURRENT* or to a file *FILE\_CONCURRENT*. Tasks that share a "CONCURRENT" parameter will be executed in parallel, if any other dependency prevents this. The *CONCURRENT* direction allows users to have access from multiple tasks to the same object/-file during their executions. However, note that COMPSs does not manage the interaction with the objects or files used/modified concurrently. Taking care of the access/modification of the concurrent objects is responsibility of the developer.

```

from pycompss.api.task import task      # Import @task decorator
from pycompss.api.parameter import *    # Import parameter metadata for the @task decorator

@task (f=FILE_CONCURRENT)
def func(f, i):
    ...

```

Moreover, it is possible to specify that a parameter is a collection of elements (e.g. list) and its direction (*COLLECTION\_IN* or *COLLECTION\_INOUT*). In this case, the list may contain sub-objects that will be handled automatically by the runtime. It is important to annotate data structures as collections if in other tasks there are accesses to individual elements of these collections as parameters. Without this annotation, the runtime will not be able to identify data dependences between the collections and the individual elements.

```

from pycompss.api.task import task      # Import @task decorator
from pycompss.api.parameter import COLLECTION_IN # Import parameter metadata for the @task decorator

@task (my_collection=COLLECTION_IN)
def func(my_collection):
    for element in my_collection:
        ...

```

The sub-objects of the collection can be collections of elements (and recursively). In this case, the runtime also keeps track of all elements contained in all sub-collections. In order to improve the performance, the depth of the sub-objects can be limited through the use of the *depth* parameter as:

```

@task (my_collection={Type: COLLECTION_IN, Depth: 2})
def func(my_collection):
    for inner_collection in my_collection:
        for element in inner_collection:
            # The contents of element will not be tracked
            ...

```

If the function or method returns a value, the programmer must use the *returns* argument within the `@task` decorator. In this argument, the programmer can specify the type of that value:

```
@task(returns =int)
def ret_func():
    return 1
```

Moreover, if the function or method returns more than one value, the programmer can specify how many and their type in the *returns* argument. The next code snippet shows how to specify that two values (an integer and a list) are returned:

```
@task(returns =(int, list))
def ret_func():
    return 1, [2, 3]
```

Alternatively, the user can specify the number of return statements as an integer value. This way of specifying the amount of return eases the *returns* definition since the user does not need to specify explicitly the type of the return arguments. However, it must be considered that the type of the object returned when the task is invoked will be a future object. This consideration may lead to an error if the user expects to invoke a task defined within an object returned by a previous task. In this scenario, the solution is to specify explicitly the return type.

```
@task(returns =1)
def ret_func():
    return "my_string"

@task(returns =2)
def ret_func():
    return 1, [2, 3]
```

The use of *\*args* and *\*\*kwargs* as function parameters is also supported:

```
@task(returns =int)
def argkwarg_func(*args, **kwargs):
    return sum(args) + len(kwargs)
```

And even with other parameters, such as usual parameters and *default defined arguments*. The next snippet shows an example of a task with two three parameters (whose one of them ('s') has a default value), *\*args* and *\*\*kwargs*.

```
@task(returns =int)
def multiarguments_func(v, w, s = 2, *args, **kwargs):
    return (v * w) + sum(args) + len(kwargs) + s
```

For tasks corresponding to instance methods, by default the task is assumed to modify the callee object (the object on which the method is invoked). The programmer can tell

otherwise by setting the *target\_direction* argument of the *@task* decorator to *IN*.

```
class MyClass(object):
    ...
    @task(target_direction = IN)
    def instance_method(self):
        ... # self is NOT modified here
```

**NOTE:** In order to avoid serialization issues, the classes must not be declared in the same file that contains the main method (*if \_\_name\_\_ == '\_\_main\_\_':*).

### 3.1.1.1 Scheduler hints

The programmer can provide hints to the scheduler through specific arguments within the *@task* decorator.

For instance, the programmer can mark a task as a high-priority task with the *priority* argument of the *@task* decorator. In this way, when the task is free of dependencies, it will be scheduled before any of the available low-priority (regular) tasks. This functionality is useful for tasks that are in the critical path of the application's task dependency graph.

```
@task(priority =True)
def func():
    ...
```

Moreover, the user can also mark a task as distributed with the *is\_distributed* argument or as replicated with the *is\_replicated* argument. When a task is marked with *is\_distributed=True*, the method must be scheduled in a forced round robin among the available resources. On the other hand, when a task is marked with *is\_replicated=True*, the method must be executed in all the worker nodes when invoked from the main application. The default value for these parameters is *False*.

```
@task(is_distributed =True)
def func():
    ...

@task(is_replicated =True)
def func2():
    ...
```

In case a task fails, the whole application behaviour can be defined using the *on\_failure* argument. It has four possible values: 'RETRY', 'CANCEL\_SUCCESSORS', 'FAIL' and 'IGNORE'. 'RETRY' is the default behaviour, making the task to be executed again, on the same worker or in another worker if the failure remains. 'CANCEL\_SUCCESSORS' ignores the failed task and cancels the execution of the successor tasks, 'FAIL' stops the whole execution once a task fails and 'IGNORE' ignores the failure and continues with the normal execution.

```
@task(on_failure ='CANCEL_SUCCESSORS')
def func():
    ...
```

Table 1 summarizes the arguments that can be found in the `@task` decorator.

Argument	Value
Formal parameter name	<p><b>(default: empty)</b>: The parameter is an object or a simple tipe that will be inferred.</p> <ul style="list-style-type: none"> <li>- IN: read-only parameter, all types.</li> <li>- INOUT: read-write parameter, all types except file (primitives, strings, objects).</li> <li>- OUT: write-only parameter, all types except file (primitives, strings, objects).</li> <li>- CONCURRENT : concurrent read-write parameter, all types except file (primitives, strings, objects).</li> <li>- FILE/FILE_IN: read-only file parameter.</li> <li>- FILE_INOUT: read-write file parameter.</li> <li>- FILE_OUT: write-only file parameter.</li> <li>- FILE_CONCURRENT: concurrent read-write file parameter.</li> <li>- COLLECTION_IN: read-only collection parameter (list).</li> <li>- COLLECTION_INOUT: read-write collection parameter (list).</li> <li>- Dictionary: {Type:(empty=object)/FILE/COLLECTION, Direction:(empty=IN)/IN/INOUT/OUT/CONCURRENT}</li> </ul>
returns	int (for integer and boolean), long, float, str, dict, list, tuple, user-defined classes
target_direction	INOUT (default), IN or CONCURRENT
priority	True or False (default)
is_distributed	True or False (default)
is_replicated	True or False (default)
on_failure	'RETRY' (default), 'CANCEL_SUCCESSORS', 'FAIL' or 'IGNORE'

Table 1: Arguments of the `@task` decorator.

### 3.1.1.2 Other task types

In addition to this API functions, the programmer can use a set of decorators for other purposes.

For instance, there is a set of decorators that can be placed over the `@task` decorator in order to define the task methods as a binary invocation (with the `@binary` decorator),

as a OmpSs invocation (with the `@ompss` decorator), as a MPI invocation (with the `@mpi` decorator), as a COMPSs application (with the `@compss` decorator), or as a task that requires multiple nodes (with the `@multinode` decorator). These decorators must be placed over the `@task` decorator, and under the `@constraint` decorator if defined.

Consequently, the task body will be empty and the function parameters will be used as invocation parameters with some extra information that can be provided within the `@task` decorator.

The following subparagraphs describe their usage.

**3.1.1.2.1 Binary decorator** The `@binary` decorator shall be used to define that a task is going to invoke a binary executable.

In this context, the `@task` decorator parameters will be used as the binary invocation parameters (following their order in the function definition). Since the invocation parameters can be of different nature, information on their type can be provided through the `@task` decorator.

The following snippet shows the most simple binary task definition without/with constraints (without parameters):

```
from pycompss.api.task import task
from pycompss.api.binary import binary

@binary(binary="mybinary.bin")
@task()
def binary_func():
    pass

@constraint(computingUnits="2")
@binary(binary="otherbinary.bin")
@task()
def binary_func2():
    pass
```

The invocation of these tasks would be equivalent to:

```
./mybinary.bin
./otherbinary.bin # in resources that respect the constraint.
```

The following snippet shows a more complex binary invocation, with files as parameters:

```
from pycompss.api.task import task
from pycompss.api.binary import binary
from pycompss.api.parameter import *

@binary(binary="grep", working_dir=".")
@task(infile={Type:FILE_IN_STDIN}, result={Type:FILE_OUT_STDOUT})
def grepper():
    pass

# This task definition is equivalent to the following, which is more verbose:

@binary(binary="grep", working_dir=".")
```

```
@task(infile={Type:FILE_IN, Stream:STDIN}, result={Type:FILE_OUT, Stream:STDOUT})
def grepper(keyword, infile, result):
    pass

if __name__=='__main__':
    infile = "infile.txt"
    outfile = "outfile.txt"
    grepper("Hi", infile, outfile)
```

The invocation of the *grepper* task would be equivalent to:

```
# grep keyword < infile > result
grep Hi < infile.txt > outfile.txt
```

Please note that the *keyword* parameter is a string, and it is respected as is in the invocation call.

Thus, PyCOMPSs can also deal with prefixes for the given parameters. The following snippet performs a system call (*ls*) with specific prefixes:

```
from pycompss.api.task import task
from pycompss.api.binary import binary
from pycompss.api.parameter import *

@binary(binary="ls")
@task(hide={Type:FILE_IN, Prefix:"--hide="}, sort={Prefix:"--sort="})
def myLs(flag, hide, sort):
    pass

if __name__=='__main__':
    flag = '-l'
    hideFile = "fileToHide.txt"
    sort = "time"
    myLs(flag, hideFile, sort)
```

The invocation of the *myLs* task would be equivalent to:

```
# ls -l --hide=hide --sort=sort
ls -l --hide=fileToHide.txt --sort=time
```

This particular case is intended to show all the power of the *@binary* decorator in conjunction with the *@task* decorator. Please note that although the *hide* parameter is used as a prefix for the binary invocation, the *fileToHide.txt* would also be transferred to the worker (if necessary) since its type is defined as *FILE\_IN*. This feature enables to build more complex binary invocations.

**3.1.1.2.2 OmpSs decorator** The *@omps* decorator shall be used to define that a task is going to invoke a OmpSs executable.



```

from pycompss.api.ompss import ompss

@ompss(binary="ompssApp.bin")
@task()
def ompss_func():
    pass

```

The OmpSs executable invocation can also be enriched with parameters, files and prefixes as with the `@binary` decorator through the function parameters and `@task` decorator information. Please, check subparagraph 3.1.1.2.1 for more details.

**3.1.1.2.3 MPI decorator** The `@mpi` decorator shall be used to define that a task is going to invoke a MPI executable.

```

from pycompss.api.mpi import mpi

@mpi(binary="mpiApp.bin", runner="mpirun", computing_nodes=2)
@task()
def mpi_func():
    pass

```

The MPI executable invocation can also be enriched with parameters, files and prefixes as with the `@binary` decorator through the function parameters and `@task` decorator information. Please, check subparagraph 3.1.1.2.1 for more details.

**3.1.1.2.4 COMPSs decorator** The `@compss` decorator shall be used to define that a task is going to be a COMPSs application. It enables to have nested PyCOMPSs/COMPSs applications.

```

from pycompss.api.compss import compss

@compss(runcompss="{RUNCOMPSS}", flags="-d",
        app_name="/path/to/simple_compss_nested.py", computing_nodes="2")
@task()
def compss_func():
    pass

```

The COMPSs application invocation can also be enriched with the flags accepted by the `runcompss` executable. Please, check execution manual for more details about the supported flags.

**3.1.1.2.5 Multinode decorator** The `@multinode` decorator shall be used to define that a task is going to use multiple nodes (e.g. using internal parallelism).

```

from pycompss.api.multinode import multinode

@multinode(computing_nodes="2")
@task()
def multinode_func():
    pass

```

The only supported parameter is *computing\_nodes*, used to define the number of nodes required by the task (the default value is 1). The mechanism to get the number of nodes, threads and their names to the task is through the *COMPSS\_NUM\_NODES*, *COMPSS\_NUM\_THREADS* and *COMPSS\_HOSTNAMES* environment variables respectively, which are exported within the task scope by the COMPSs runtime before the task execution.

**3.1.1.2.6 Parameters summary** Table 2 summarizes the parameters of these decorators.

Decorator	Parameters
@binary	<p><b>binary:</b> (Mandatory) String defining the full path of the binary that must be executed.</p> <p><b>working_dir:</b> Full path of the binary working directory inside the COMPSs Worker.</p>
@ompss	<p><b>binary:</b> (Mandatory) String defining the full path of the binary that must be executed.</p> <p><b>working_dir:</b> Full path of the binary working directory inside the COMPSs Worker.</p>
@mpi	<p><b>binary:</b> (Mandatory) String defining the full path of the binary that must be executed.</p> <p><b>working_dir:</b> Full path of the binary working directory inside the COMPSs Worker.</p> <p><b>runner:</b> (Mandatory) String defining the MPI runner command.</p> <p><b>computing_nodes:</b> Integer defining the number of computing nodes reserved for the MPI execution (only a single node is reserved by default).</p>
@compss	<p><b>runcompss:</b> (Mandatory) String defining the full path of the runcompss binary that must be executed.</p> <p><b>flags:</b> String defining the flags needed for the runcompss execution.</p> <p><b>app_name:</b> (Mandatory) String defining the application that must be executed.</p> <p><b>computing_nodes:</b> Integer defining the number of computing nodes reserved for the COMPSs execution (only a single node is reserved by default).</p>
@multinode	<p><b>computing_nodes:</b> Integer defining the number of computing nodes reserved for the task execution (only a single node is reserved by default).</p>

Table 2: @binary, @ompss, @mpi, @compss and @multinode decorators supported parameters.

In addition to the parameters that can be used within the `@task` decorator described in Table 1, Table 3 summarizes the *Stream* parameter that can be used within the `@task` decorator for the function parameters when using the `@binary`, `@ompss` and `@mpi` decorators. In particular, the *Stream* parameter is used to indicate that a parameter is going to be considered as a *FILE* but as a stream (e.g. `>`, `<` and `2 >` in bash) for the `@binary`, `@ompss` and `@mpi` calls.

Parameter key	Values
<b>Stream</b>	<p><b>(default: empty):</b> Not a stream.</p> <p><b>STDIN:</b> Standard input.</p> <p><b>STDOUT:</b> Standard output.</p> <p><b>STDERR:</b> Standard error.</p>

Table 3: Supported Streams for the `@binary`, `@ompss` and `@mpi` decorators

Moreover, there are some shortcuts that can be used for files type definition as parameters within the `@task` decorator (Table 4). It is not necessary to indicate the *Direction* nor the *Stream* since it may be already be indicated with the shortcut.

Alias	Description
<b>COLLECTION(_IN)</b>	Type: COLLECTION, Direction: IN
<b>COLLECTION_INOUT</b>	Type: COLLECTION, Direction: INOUT
<b>FILE(_IN)_STDIN</b>	Type: File, Direction: IN, Stream: STDIN
<b>FILE(_IN)_STDERR</b>	Type: File, Direction: IN, Stream: STDOUT
<b>FILE(_IN)_STDOUT</b>	Type: File, Direction: IN, Stream: STDERR
<b>FILE_OUT_STDIN</b>	Type: File, Direction: OUT, Stream: STDIN
<b>FILE_OUT_STDERR</b>	Type: File, Direction: OUT, Stream: STDOUT
<b>FILE_OUT_STDOUT</b>	Type: File, Direction: OUT, Stream: STDERR
<b>FILE_INOUT_STDIN</b>	Type: File, Direction: INOUT, Stream: STDIN
<b>FILE_INOUT_STDERR</b>	Type: File, Direction: INOUT, Stream: STDOUT
<b>FILE_INOUT_STDOUT</b>	Type: File, Direction: INOUT, Stream: STDERR
<b>FILE_CONCURRENT</b>	Type: File, Direction: CONCURRENT
<b>FILE_CONCURRENT_STDIN</b>	Type: File, Direction: CONCURRENT, Stream: STDIN

<b>FILE_CONCURRENT_STDERR</b>	Type: File, Direction: CONCURRENT, Stream: STDOUT
<b>FILE_CONCURRENT_STDOUT</b>	Type: File, Direction: CONCURRENT, Stream: STDERR

Table 4: File parameters definition shortcuts

These parameter keys, as well as the shortcuts, can be imported from the PyCOMPSs library:

```
from pycompss.api.parameter import *
```

### 3.1.2 Constraints

As in Java COMPSs applications, it is possible to define constraints for each task. To this end, the decorator `@constraint` followed by the desired constraints needs to be placed over the `@task` decorator.

```
from pycompss.api.task import task
from pycompss.api.constraint import constraint
from pycompss.api.parameter import INOUT

@constraint (computing_units="4")
@task (c=INOUT)
def func(a, b, c):
    c += a*b
    ...
```

This decorator enables the user to set the particular constraints for each task, such as the amount of Cores required explicitly. Alternatively, it is also possible to indicate that the value of a constraint is specified in a environment variable. A full description of the supported constraints can be found in Table 8 in Section 5.

For example:

```
from pycompss.api.task import task
from pycompss.api.constraint import constraint
from pycompss.api.parameter import INOUT

@constraint (computing_units="4",
            app_software="numpy,scipy,gnuplot",
            memory_size="$MIN_MEM_REQ")
@task (c=INOUT)
def func(a, b, c):
    c += a*b
    ...
```

Or another example requesting a CPU core and a GPU:

```

from pycompss.api.task import task
from pycompss.api.constraint import constraint

@constraint (processors=[{'processorType':'CPU', 'computingUnits':'1'},
                        {'processorType':'GPU', 'computingUnits':'1'}])
@task (returns=1)
def func(a, b, c):
    ...
    return result

```

When the task requests a GPU, COMPSs provides the information about the assigned GPU through the *COMPSS\_BINDED\_GPUS*, *CUDA\_VISIBLE\_DEVICES* and *GPU\_DEVICE\_ORDINAL* environment variables. This information can be gathered from the task code in order to use the GPU.

Please, take into account that in order to respect the constraints, the peculiarities of the infrastructure must be defined in the *resources.xml* file.

### 3.1.3 Implements

As in Java COMPSs applications, it is possible to define multiple implementations for each task. In particular, a programmer can define a task for a particular purpose, and multiple implementations for that task with the same objective, but with different constraints (e.g. specific libraries, hardware, etc). To this end, the *@implement* decorator followed with the specific implementations constraints (with the *@constraint* decorator, see Section 3.1.2) needs to be placed over the *@task* decorator. Although the user only calls the task that is not decorated with the *@implement* decorator, when the application is executed in a heterogeneous distributed environment, the runtime will take into account the constraints on each implementation and will try to invoke the implementation that fulfills the constraints within each resource, keeping this management invisible to the user.

```

from pycompss.api.implement import implement

@implement(source_class="sourcemodule", method="main_func")
@constraint(app_software="numpy")
@task(returns=list)
def myfunctionWithNumpy(list1, list2):
    # Operate with the lists using numpy
    return resultList

@task(returns=list)
def main_func(list1, list2):
    # Operate with the lists using built-in functions
    return resultList

```

Please, note that if the implementation is used to define a binary, OmpSs, MPI, COMPSs or multinode task invocation (see Section 3.1.1.2), the *@implement* decorator must be always on top of the decorators stack, followed by the *@constraint* decorator, then the *@binary/@ompss/@mpi/@compss/@multinode* decorator, and finally, the *@task* decorator in the lowest level.

### 3.1.4 Main Program

The main program of the application is a sequential code that contains calls to the selected tasks. In addition, when synchronizing for task data from the main program, there exist four API functions that can be invoked:

- `compss_open(file_name, mode='r')`: similar to the Python `open()` call. It synchronizes for the last version of file `file_name` and returns the file descriptor for that synchronized file. It can have an optional parameter `mode`, which defaults to `'r'`, containing the mode in which the file will be opened (the open modes are analogous to those of Python `open()`).
- `compss_delete_file(file_name)`: notifies the runtime to delete a file.
- `compss_wait_on_file(file_name)`: synchronizes for the last version of the file `file_name`. and returns `True` if success (`False` otherwise).
- `compss_delete_object(object)`: notifies the runtime to delete all the associated files to a given object.
- `compss_barrier(no_more_tasks=False)`: performs a explicit synchronization, but does not return any object. The use of `compss_barrier()` forces to wait for all tasks that have been submitted before the `compss_barrier()` is called. When all tasks submitted before the `compss_barrier()` have finished, the execution continues. The `no_more_tasks` is used to specify if no more tasks are going to be submitted after the `compss_barrier()`.
- `compss_wait_on(obj, to_write=True)`: synchronizes for the last version of object `obj` and returns the synchronized object. It can have an optional boolean parameter `to_write`, which defaults to `True`, that indicates whether the main program will modify the returned object. It is possible to wait on a list of objects. In this particular case, it will synchronize all future objects contained in the list.

To illustrate the use of the aforementioned API functions, the following example first invokes a task `func` that writes a file, which is later synchronized by calling `compss_open()`. Later in the program, an object of class `MyClass` is created and a task method `method` that modifies the object is invoked on it; the object is then synchronized with `compss_wait_on()`, so that it can be used in the main program from that point on.

Then, a loop calls again ten times to `func` task. Afterwards, the barrier performs a synchronization, and the execution of the main user code will not continue until the ten `func` tasks have finished.

```
from pycompss.api.api import compss_open
from pycompss.api.api import compss_delete_file
from pycompss.api.api import compss_wait_on
from pycompss.api.api import compss_barrier

if __name__=='__main__':
    my_file = 'file.txt'
    func(my_file)
    fd = compss_open(my_file)
    ...
```

```

my_file2 = 'file2.txt'
func(my_file2)
fd = compss_delete_file(my_file2)
...

my_obj = MyClass()
my_obj.method()
my_obj = compss_wait_on(my_obj)
...

for i in range(10):
    func(str(i) + my_file)
compss_barrier()
...

```

The corresponding task selection for the example above would be:

```

@task(f=FILE_OUT)
def func(f):
    ...

class MyClass(object):
    ...

    @task()
    def method(self):
        ... # self is modified here

```

Table 5 summarizes the API functions to be used in the main program of a COMPSs Python application.

Function	Use
<code>compss_open(file_name, mode='r')</code>	Synchronizes for the last version of a file and returns its file descriptor.
<code>compss_delete_file(file_name)</code>	Notifies the runtime to remove a file.
<code>compss_wait_on_file(file_name)</code>	Synchronizes for the last version of a file.
<code>compss_delete_object(object)</code>	Notifies the runtime to delete the associated file to this object.
<code>compss_barrier(no_more_tasks=False)</code>	Wait for all tasks submitted before the barrier.
<code>compss_wait_on(obj, to_write=True)</code>	Synchronizes for the last version of an object (or a list of objects) and returns it.

Table 5: COMPSs Python API functions.

Besides the synchronization API functions, the programmer has also a decorator for automatic function parameters synchronization at his disposal. The `@local` decorator can be placed over functions that are not decorated as tasks, but that may receive results

from tasks. In this case, the `@local` decorator synchronizes the necessary parameters in order to continue with the function execution without the need of using explicitly the `comps_wait_on` call for each parameter.

```
from pycomps.api.task import task
from pycomps.api.api import comps_wait_on
from pycomps.api.parameter import INOUT
from pycomps.api.local import local

@task(returns =list)
@task(v=INOUT)
def append_three_ones(v):
    v += [1, 1, 1]

@local
def scale_vector(v, k):
    return [k*x for x in v]

if __name__=='__main__':
    v = [1,2,3]
    append_three_ones(v)
    # v is automatically synchronized when calling the scale_vector function.
    w = scale_vector(v, 2)
```

### 3.1.5 Important Notes

If the programmer selects as a task a function or method that returns a value, that value is not generated until the task executes.

```
@task(returns =MyClass)
def ret_func():
    return MyClass(...)

...

if __name__=='__main__':
    # o is a future object
    o = ret_func()
```

The object returned can be involved in a subsequent task call, and the COMPSs runtime will automatically find the corresponding data dependency. In the following example, the object `o` is passed as a parameter and callee of two subsequent (asynchronous) tasks, respectively:

```
if __name__=='__main__':
    # o is a future object
    o = ret_func()

    ...

    another_task(o)

    ...

    o.yet_another_task()
```



In order to synchronize the object from the main program, the programmer has to synchronize (using the *compss\_wait\_on* function) in the same way as with any object updated by a task:

```
if __name__=='__main__':
    # o is a future object
    o = ret_func()

    ...

    o = compss_wait_on(o)
```

Moreover, it is possible to synchronize a list of objects. This is particularly useful when the programmer expect to synchronize more than one elements (using the *compss\_wait\_on* function):

```
if __name__=='__main__':
    # l is a list of objects where some/all of them may be future objects
    l = []
    for i in range(10):
        l.append(ret_func())

    ...

    l = compss_wait_on(l)
```

For instances of user-defined classes, the classes of these objects should have an empty constructor, otherwise the programmer will not be able to invoke task instance methods on those objects:

```
# In file utils.py
class MyClass(object):
    def __init__(self): # empty constructor
        ...

    ...

# In file main.py
from utils import MyClass

if __name__=='__main__':
    o = ret_func()

    # invoking a task instance method on a future object can only
    # be done when an empty constructor is defined in the object's
    # class
    o.yet_another_task()
```

In order to make the COMPSs Python binding function correctly, the programmer should not use relative imports in the code. Relative imports can lead to ambiguous code and they are discouraged in Python, as explained in:

<http://docs.python.org/2/faq/programming.html#what-are-the-best-practices-for-using-import-in-a-module>

## 3.2 Application Execution

The next subsections describe how to execute applications with the COMPSs Python binding.

### 3.2.1 Environment

The following environment variables must be defined before executing a COMPSs Python application:

JAVA\_HOME: Java JDK installation directory (e.g. `/usr/lib/jvm/java-8-openjdk/`)

### 3.2.2 Command

In order to run a Python application with COMPSs, the `runcompss` script can be used, like for Java and C/C++ applications. An example of an invocation of the script is:

```
compss@bsc:~$ runcompss \  
  --lang=python \  
  --pythonpath=$TEST_DIR \  
  --library_path=/home/user/libdir \  
  $TEST_DIR/test.py arg1 arg2
```

Or alternatively, use the `pycompss` module:

```
compss@bsc:~$ python -m pycompss \  
  --pythonpath=$TEST_DIR \  
  --library_path=/home/user/libdir \  
  $TEST_DIR/test.py arg1 arg2
```

For full description about the options available for the `runcompss` command please check the *COMPSs User Manual: Application Execution* available at <http://compss.bsc.es>.

## 3.3 Development with Jupyter notebook

PyCOMPSs can also be used within Jupyter notebooks. This feature allows users to develop and run their PyCOMPSs applications in a Jupyter notebook, where it is possible to modify the code during the execution and experience an interactive behaviour.

### 3.3.1 Environment

The following libraries must be present in the appropriate environment variables in order to enable PyCOMPSs within Jupyter notebook:

PYTHONPATH: The path where PyCOMPSs is installed (e.g. `/opt/COMPSs/Bindings/python/`)

LD\_LIBRARY\_PATH: The path where the `libbindings-commons.so` library is located (e.g. `/opt/COMPSs/Bindings/bindings-common/lib/`) and the path where the `libjvm.so` library is located (e.g. `/usr/lib/jvm/java-8-openjdk/jre/lib/amd64/server/`).

### 3.3.2 API calls

In this case, the user is responsible of starting and stopping the COMPSs runtime. To this end, PyCOMPSs provides a module that provides two API calls: one for starting the COMPSs runtime, and another for stopping it.

This module can be imported from the *pycompss* library:

```
import pycompss.interactive as ipycompss
```

And contains two main functions: *start* and *stop*. These functions can then be invoked as follows for the COMPSs runtime deployment with default parameters:

```
# Previous user code
ipycompss.start()

# User code that can benefit from PyCOMPSs

ipycompss.stop()

# Subsequent code
```

Between the *start* and *stop* function calls, the user can write its own python code including PyCOMPSs imports, decorators and synchronization calls described in Section 3.1. The code can be splitted into multiple cells.

The *start* and *stop* functions accept parameters in order to customize the COMPSs runtime (such as the flags that can be selected with the “runcompss” command). Table 6 summarizes the accepted parameters of the *start* function. Table 7 summarizes the accepted parameters of the *stop* function.

Parameter Name	Parameter Type	Description
log_level	String	PyCOMPSs log level. Options: “off”, “info” and “debug”. (Default: “off”)
debug	Boolean	COMPSs runtime debug (Default: False) (overrides log_level)
o_c	Boolean	Object conversion to string when possible (Default: False)
graph	Boolean	Task dependency graph generation (Default: False)
trace	Boolean	Paraver trace generation (Default: False)
monitor	Integer	Monitor refresh rate (Default: None - Monitoring disabled)
project_xml	String	Path to the project XML file (Default: \$COMPSS/Runtime/configuration/xml/projects/default_project.xml)
resources_xml	String	Path to the resources XML file (Default: \$COMPSS/Runtime/configuration/xml/resources/default_resources.xml)

summary	Boolean	Show summary at the end of the execution (Default: False)
storage_impl	String	Path to an storage implementation (Default: None)
storage_conf	String	Storage configuration file path (Default: None)
task_count	Integer	Number of task definitions (Default: 50)
app_name	String	Application name (Default: "Interactive")
uuid	String	Application uuid (Default: None - Will be random)
base_log_dir	String	Base directory to store COMPSs log files (a .COMPSs/ folder will be created inside this location) (Default: User homeBase log path)
specific_log_dir	String	Use a specific directory to store COMPSs log files (the folder MUST exist and no sandbox is created) (Default: Disabled)
extrae_cfg	String	Sets a custom extrae config file. Must be in a shared disk between all COMPSs workers (Default: None)
comm	String	Class that implements the adaptor for communications. Supported adaptors: "es.bsc.compss.nio.master.NIOAdaptor" and "es.bsc.compss.gat.master.GATAaptor" (Default: "es.bsc.compss.nio.master.NIOAdaptor")
conn	String	Class that implements the runtime connector for the cloud. Supported connectors: "es.bsc.compss.connectors.DefaultSSHConnector" and "es.bsc.compss.connectors.DefaultNoSSHConnector" (Default: es.bsc.compss.connectors.DefaultSSHConnector)
master_name	String	Hostname of the node to run the COMPSs master (Default: "")
master_port	String	Port to run the COMPSs master communications. Only for NIO adaptor (Default: "[43000,44000]")
scheduler	String	Class that implements the Scheduler for COMPSs. Supported schedulers: "es.bsc.compss.scheduler.fullGraphScheduler.FullGraphScheduler", "es.bsc.compss.scheduler.fifoScheduler.FIFOScheduler" and "es.bsc.compss.scheduler.resourceEmptyScheduler.ResourceEmptyScheduler" (Default: "es.bsc.compss.scheduler.loadBalancingScheduler.LoadBalancingScheduler")
jvm_workers	String	Extra options for the COMPSs Workers JVMs. Each option separated by "," and without blank spaces (Default: "-Xms1024m,-Xmx1024m,-Xmn400m")
cpu_affinity	String	Sets the CPU affinity for the workers. Supported options: "disabled", "automatic", user defined map of the form "0-8/9,10,11/12-14,15,16" (Default: "automatic")

gpu_affinity	String	Sets the GPU affinity for the workers. Supported options: “disabled”, “automatic”, user defined map of the form “0-8/9,10,11/12-14,15,16” (Default: “automatic”)
profile_input	String	Path to the file which stores the input application profile (Default: “”)
profile_output	String	Path to the file to store the application profile at the end of the execution (Default: “”)
scheduler_config	String	Path to the file which contains the scheduler configuration (Default: “”)
external_adaptation	Boolean	Enable external adaptation. This option will disable the Resource Optimizer (Default: False)
propagate_virtual_environment	Boolean	Propagate the master virtual environment to the workers (Default: False)
verbose	Boolean	Verbose mode (Default: False)

Table 6: PyCOMPSs start function for Jupyter notebook

Parameter Name	Parameter Type	Description
sync	Boolean	Synchronize the objects left on the user scope. (Default: False)

Table 7: PyCOMPSs stop function for Jupyter notebook

The following code snippet shows how to start a COMPSs runtime with tracing and graph generation enabled (with *trace* and *graph* parameters), as well as enabling the monitor with a refresh rate of 2 seconds (with the *monitor* parameter). It also synchronizes all remaining objects in the scope with the *sync* parameter when invoking the *stop* function.

```
# Previous user code
ipycompss.start(graph=True, trace=True, monitor=2000)

# User code that can benefit from PyCOMPSs

ipycompss.stop(sync=True)

# Subsequent code
```

### 3.3.3 Application execution

The application can be executed as a common Jupyter notebook by steps or the whole application.

Please, note that once the COMPSs runtime has been stopped it is necessary to restart the python kernel in Jupyter before starting another COMPSs runtime. To this end, click on “Kernel” and “Restart” (or “Restart & Clear Output” or “Restart & Run All”, depending on the need).

## 3.4 Integration with Numba

PyCOMPSs can also be used with Numba. Numba (<http://numba.pydata.org/>) is an Open Source JIT compiler for Python which provides a set of decorators and functionalities to translate Python functions to optimized machine code.

### 3.4.1 Basic usage

PyCOMPSs’ tasks can be decorated with Numba’s `@jit/@njit` decorator (with the appropriate parameters) just below the `@task` decorator in order to apply Numba to that task.

```
from pycomps.api.task import task      # Import @task decorator
from numba import jit

@task (returns=1)
@jit ()
def numba_func(a, b):
    ...
```

### 3.4.2 Advanced usage

PyCOMPSs can be also used in conjunction with the Numba’s `@vectorize`, `@guvectorize`, `@stencil` and `@cfunc`. But since these decorators do not preserve the original argument specification of the original function, their usage is done through the `numba` parameter within the `@task` decorator. This parameter accepts:

- Boolean: True: Applies *jit* to the function.
- Dictionary{k, v}: Applies *jit* with the dictionary parameters to the function (allows to specify specific jit parameters (e.g. `nopython=True`)).
- String:
  - “jit”: Applies *jit* to the function.
  - “njit”: Applies *jit* with `nopython=True` to the function.
  - “generated\_jit”: Applies *generated\_jit* to the function.
  - “vectorize”: Applies *vectorize* to the function.Needs some extra flags in the `@task` decorator:

- *numba.signature*: String with the *vectorize* signature.

“guvectorize”: Applies *guvectorize* to the function.

Needs some extra flags in the *@task* decorator:

- *numba.signature*: String with the *guvectorize* signature.

- *numba.declaration*: String with the *guvectorize* declaration.

“stencil”: Applies *stencil* to the function.

“cfunc”: Applies *cfunc* to the function.

Needs some extra flags in the *@task* decorator:

- *numba.signature*: String with the *cfunc* signature.

Moreover, the *@task* decorator also allows to define specific flags for the *jit*, *njit*, *generated\_jit*, *vectorize*, *guvectorize* and *cfunc* functionalities with the *numba\_flags* hint. This hint is used to declare a dictionary with the flags expected to use with these numba functionalities. The default flag included by PyCOMPSs is the *cache=True* in order to exploit the function caching of Numba across tasks.

For example, to apply *jit* to a function:

```
from pycompss.api.task import task

@task (numba='jit') # Aternatively: @task(numba=True)
def jit_func(a, b):
    ...
```

And if the developer wants to use specific flags with *jit* (e.g. *parallel=True*):

```
from pycompss.api.task import task

@task (numba='jit', numba_flags={'parallel':True})
def jit_func(a, b):
    ...
```

Other Numba’s functionalities require the specification of the function signature and declaration. In the next example a task that will use the *vectorize* with three parameters and a specific flag to target the cpu is shown:

```
from pycompss.api.task import task

@task (returns=1,
      numba='vectorize',
      numba_signature=['float32(float32, float32, float32)'],
      numba_flags={'target':'cpu'})
def vectorize_task(a, b, c):
    return a * b * c
```

Details about numba and the specification of the signature, declaration and flags can be found in the Numba’s webpage (<http://numba.pydata.org/>).

## 4 C/C++ Binding

COMPSs provides a binding for C and C++ applications. The new C++ version in the current release comes with support for objects as task parameters and the use of class methods as tasks.

### 4.1 Programming Model

#### 4.1.1 Task Selection

As in Java the user has to provide a task selection by means of an interface. In this case the interface file has the same name as the main application file plus the suffix “idl”, i.e. Matmul.idl, where the main file is called Matmul.cc.

```
interface Matmul
{
    // C functions
    void initMatrix(inout Matrix matrix,
                   in int mSize,
                   in int nSize,
                   in double val);

    void multiplyBlocks(inout Block block1,
                       inout Block block2,
                       inout Block block3);
};
```

The syntax of the interface file is shown in the previous code. Tasks can be declared as classic C function prototypes, this allow to keep the compatibility with standard C applications. In the example, `initMatrix` and `multiplyBlocks` are functions declared using its prototype, like in a C header file, but this code is C++ as they have objects as parameters (objects of type `Matrix`, or `Block`).

The grammar for the interface file is:

```
["static"] return-type task-name ( parameter {, parameter }* );
return-type = "void" | type
ask-name = <qualified name of the function or method>
parameter = direction type parameter-name
direction = "in" | "out" | "inout"
type = "char" | "int" | "short" | "long" | "float" | "double" | "boolean" |
      "char[<size>]" | "int[<size>]" | "short[<size>]" | "long[<size>]" |
      "float[<size>]" | "double[<size>]" | "string" | "File" | class-name
class-name = <qualified name of the class>
```

#### 4.1.2 Main Program

The next listing includes an example of matrix multiplication written in C++.



```

#include "Matmul.h"
#include "Matrix.h"
#include "Block.h"
int N; //MSIZE
int M; //BSIZE
double val;
int main(int argc, char **argv)
{
    Matrix A;
    Matrix B;
    Matrix C;

    N = atoi(argv[1]);
    M = atoi(argv[2]);
    val = atof(argv[3]);

    compss_on();

    A = Matrix::init(N,M,val);

    initMatrix(&B,N,M,val);
    initMatrix(&C,N,M,0.0);

    cout << "Waiting for initialization...\n";

    compss_wait_on(B);
    compss_wait_on(C);

    cout << "Initialization ends...\n";

    C.multiply(A, B);

    compss_off();
    return 0;
}

```

The developer has to take into account the following rules:

1. A header file with the same name as the main file must be included, in this case **Matmul.h**. This header file is automatically generated by the binding and it contains other includes and type-definitions that are required.
2. A call to the **compss\_on** binding function is required to turn on the COMPSs runtime.
3. As in C language, out or inout parameters should be passed by reference by means of the “&” operator before the parameter name.
4. Synchronization on a parameter can be done calling the **compss\_wait\_on** binding function. The argument of this function must be the variable or object we want to synchronize.
5. There is an **implicit synchronization** in the init method of Matrix. It is not possible to know the address of “A” before exiting the method call and due to this it is necessary to synchronize before for the copy of the returned value into “A” for it to be correct.
6. A call to the **compss\_off** binding function is required to turn off the COMPSs runtime.

### 4.1.3 Binding API

Besides the aforementioned **compss\_on**, **compss\_off** and **compss\_wait\_on** functions, the C/C++ main program can make use of a variety of other API calls to better manage the synchronization of data generated by tasks. These calls are as follows:

- *void compss\_ifstream(char \* filename, ifstream & ifs)*: given an uninitialized input stream *ifs* and a file *filename*, this function will synchronize the content of the file and initialize *ifs* to read from it.
- *void compss\_ofstream(char \* filename, ofstream & ofs)*: behaves the same way as *compss\_ifstream*, but in this case the opened stream is an output stream, meaning it will be used to write to the file.
- *FILE\* compss\_fopen(char \* file\_name, char \* mode)*: similar to the C/C++ *fopen* call. Synchronizes with the last version of file *file\_name* and returns the *FILE\** pointer to further reference it. As the mode parameter it takes the same that can be used in *fopen* (*r*, *w*, *a*, *r+*, *w+* and *a+*).
- *void compss\_wait\_on(T\* & obj) or T compss\_wait\_on(T & obj)*: synchronizes for the last version of object *obj*, meaning that the execution will stop until the value of *obj* up to that point of the code is received (and thus all tasks that can modify it have ended).
- *void compss\_delete\_file(char \* file\_name)*: makes an asynchronous delete of file *file\_name*. When all previous tasks have finished updating the file, it is deleted.
- *void compss\_delete\_object(T\* & obj)*: makes an asynchronous delete of an object. When all previous tasks have finished updating the object, it is deleted.
- *void compss\_barrier()*: similarly to the Python binding, performs an explicit synchronization without a return. When a *compss\_barrier* is encountered, the execution will not continue until all the tasks submitted before the *compss\_barrier* have finished.

### 4.1.4 Functions file

The implementation of the tasks in a C or C++ program has to be provided in a functions file. Its name must be the same as the main file followed by the suffix “-functions”. In our case *Matmul-functions.cc*.

```
#include "Matmul.h"
#include "Matrix.h"
#include "Block.h"

void initMatrix(Matrix *matrix, int mSize, int nSize, double val){
    *matrix = Matrix::init(mSize, nSize, val);
}

void multiplyBlocks(Block *block1, Block *block2, Block *block3){
    block1->multiply(*block2, *block3);
}
```

In the previous code, class methods have been encapsulated inside a function. This is useful when the class method returns an object or a value and we want to avoid the explicit synchronization when returning from the method.

#### 4.1.5 Additional source Files

Other source files needed by the user application must be placed under the directory “src”. In this directory the programmer must provide a **Makefile** that compiles such source files in the proper way. When the binding compiles the whole application it will enter into the src directory and execute the Makefile.

It generates two libraries, one for the master application and another for the worker application. The directive COMPSS\_MASTER or COMPSS\_WORKER must be used in order to compile the source files for each type of library. Both libraries will be copied into the lib directory where the binding will look for them when generating the master and worker applications.

#### 4.1.6 Class Serialization

In case of using an object as method parameter, as callee or as return of a call to a function, the object has to be serialized. The serialization method has to be provided inline in the header file of the object’s class by means of the “**boost**” library. The next listing contains an example of serialization for two objects of the Block class.

```
#ifndef BLOCK_H
#define BLOCK_H

#include <vector>
#include <boost/archive/text_iarchive.hpp>
#include <boost/archive/text_oarchive.hpp>
#include <boost/serialization/serialization.hpp>
#include <boost/serialization/access.hpp>
#include <boost/serialization/vector.hpp>

using namespace std;
using namespace boost;
using namespace serialization;

class Block {
public:
    Block(){};
    Block(int bSize);
    static Block *init(int bSize, double initVal);
    void multiply(Block block1, Block block2);
    void print();

private:
    int M;
    std::vector< std::vector< double > > data;

    friend class::serialization::access;
    template<class Archive>
    void serialize(Archive & ar, const unsigned int version) {
        ar & M;
        ar & data;
    }
};
#endif
```

For more information about serialization using “boost” visit the related documentation at [www.boost.org](http://www.boost.org).

#### 4.1.7 Method - Task

A task can be a C++ class method. A method can return a value, modify the *this* object, or modify a parameter.

If the method has a return value there will be an implicit synchronization before exit the method, but for the *this* object and parameters the synchronization can be done later after the method has finished.

This is because the *this* object and the parameters can be accessed inside and outside the method, but for the variable where the returned value is copied to, it can't be known inside the method.

```
#include "Block.h"

Block::Block(int bSize) {
    M = bSize;
    data.resize(M);
    for (int i=0; i<M; i++) {
        data[i].resize(M);
    }
}

Block *Block::init(int bSize, double initVal) {
    Block *block = new Block(bSize);
    for (int i=0; i<bSize; i++) {
        for (int j=0; j<bSize; j++) {
            block->data[i][j] = initVal;
        }
    }
    return block;
}

#ifdef COMPSS_WORKER

void Block::multiply(Block block1, Block block2) {
    for (int i=0; i<M; i++) {
        for (int j=0; j<M; j++) {
            for (int k=0; k<M; k++) {
                data[i][j] += block1.data[i][k] * block2.data[k][j];
            }
        }
    }
    this->print();
}

#endif

void Block::print() {
    for (int i=0; i<M; i++) {
        for (int j=0; j<M; j++) {
            cout << data[i][j] << " ";
        }
        cout << "\r\n";
    }
}
```

### 4.1.8 Task Constraints

The C/C++ binding also supports the definition of task constraints. The task definition specified in the IDL file must be decorated/annotated with the *@Constraints*. Below, you can find an example of how to define a task with a constraint of using 4 cores. The list of constraints which can be defined for a task can be found in Section 5

```
interface Matmul
{
    @Constraints(ComputingUnits = 4
    void multiplyBlocks(inout Block block1,
                       in Block block2,
                       in Block block3);
};
```

### 4.1.9 Task Versions

Another COMPSs functionality supported in the C/C++ binding is the definition of different versions for a tasks. The following code shows an IDL file where a function has two implementations, with their corresponding constraints. It shows an example where the *multiplyBlocks\_GPU* is defined as an implementation of *multiplyBlocks* using the annotation/decoration *@Implements*. It also shows how to set a processor constraint which requires a GPU processor and a CPU core for managing the offloading of the computation to the GPU.

```
interface Matmul
{
    @Constraints(ComputingUnits=4);
    void multiplyBlocks(inout Block block1,
                       in Block block2,
                       in Block block3);

    // GPU implementation
    @Constraints(processors={
        @Processor(ProcessorType=CPU, ComputingUnits=1)});
        @Processor(ProcessorType=GPU, ComputingUnits=1)});
    @Implements(multiplyBlocks);
    void multiplyBlocks_GPU(inout Block block1,
                           in Block block2,
                           in Block block3);
};
```

## 4.2 Application Compilation

To compile user's applications with the C/C++ binding two commands are used: The “**compss\_build\_app**” command allows to compile applications for a single architecture, and the “**compss\_build\_app\_multi\_arch**” command for multiple architectures. Both commands must be executed in the directory of the main application code.

### 4.2.1 Single architecture

The user command “**compss\_build\_app**” compiles both master and worker for a single architecture (e.g. x86-64, armhf, etc). Thus, whether you want to run your application in Intel based machine or ARM based machine, this command is the tool you need.

Therefore, let’s see two examples, first, the application is going to be build for the native architecture, in our case *x86-64*, and then for a target architecture, for instance *armhf*. Please note that to use cross compilation features and multiple architecture builds, you need to do the proper installation of COMPSs, find more information in the builders README.

When the target is the native architecture, the command to execute is very simple;

```
user@localhost:~/matmul_objects$ compss_build_app Matmul
[ INFO ] Java libraries are searched in the directory: /usr/lib/jvm/java-1.8.0-openjdk-amd64/jre/lib/amd64
/server
[ INFO ] Boost libraries are searched in the directory: /usr/lib/

...

[Info] The target host is: x86_64-linux-gnu

Building application for master...
g++ -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc Matrix.cc
ar rvs libmaster.a Block.o Matrix.o
ranlib libmaster.a

Building application for workers...
g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc -o Block.o
g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Matrix.cc -o Matrix.o
ar rvs libworker.a Block.o Matrix.o
ranlib libworker.a

...

Command successful.
```

In order to build an application for a different architecture e.g. *armhf*, an environment must be provided, indicating the compiler used to cross-compile, and also the location of some COMPSs dependences such as java or boost which must be compliant with the target architecture. This environment is passed by flags and arguments;

```
user@localhost:~/matmul_objects$ compss_build_app --cross-compile --cross-compile-prefix=arm-linux-
gnueabihf --java_home=/usr/lib/jvm/java-1.8.0-openjdk-armhf Matmul
[ INFO ] Java libraries are searched in the directory: /usr/lib/jvm/java-1.8.0-openjdk-armhf/jre/lib/arm/
server
[ INFO ] Boost libraries are searched in the directory: /usr/lib/
[ INFO ] You enabled cross-compile and the prefix to be used is: arm-linux-gnueabihf-

...

[ INFO ] The target host is: arm-linux-gnueabihf

Building application for master...
g++ -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc Matrix.cc
ar rvs libmaster.a Block.o Matrix.o
ranlib libmaster.a

Building application for workers...
g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc -o Block.o
g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Matrix.cc -o Matrix.o
```

```

ar rvs libworker.a Block.o Matrix.o
ranlib libworker.a

...

Command successful.

```

*[The previous outputs has been cut for simplicity]*

The **cross-compile** flag is used to indicate the users desire to cross-compile the application. It enables the use of **cross-compile-prefix** flag to define the prefix for the cross-compiler. Setting `$CROSS_COMPILE` environment variable will also work (in case you use the environment variable, the prefix passed by arguments is overridden with the variable value). This prefix is added to `$CC` and `$CXX` to be used by the user *Makefile* and lastly by the *GNU toolchain*. Regarding java and boost, **java\_home** and **boostlib** flags are used respectively. In this case, users can also use the `$JAVA_HOME` and `$BOOST_LIB` variables to indicate the java and boost for the target architecture. Note that these last arguments are purely for linkage, where `$LD_LIBRARY_PATH` is used by *Unix/Linux* systems to find libraries, so feel free to use it if you want to avoid passing some environment arguments.

## 4.2.2 Multiple architectures

The user command “**compss\_build\_app\_multi\_arch**” allows a to compile an application for several architectures. Users are able to compile both master and worker for one or more architectures. Environments for the target architectures are defined in a file specified by **cfg** flag. Imagine you wish to build your application to run the master in your Intel-based machine and the worker also in your native machine and in an ARM-based machine, without this command you would have to execute several times the command for a single architecture using its cross compile features. With the multiple architecture command is done in the following way.

```

user@localhost:~/matmul_objects$ compss_build_app_multi_arch --master=x86_64-linux-gnu --worker=arm-linux-
gnueabihf,x86_64-linux-gnu Matmul

[ INFO ] Using default configuration file: /opt/COMPSS/Bindings/c/cfgs/compssrc.
[ INFO ] Java libraries are searched in the directory: /usr/lib/jvm/java-1.8.0-openjdk-amd64/jre/lib/amd64/
server
[ INFO ] Boost libraries are searched in the directory: /usr/lib/

...

Building application for master...
g++ -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc Matrix.cc
ar rvs libmaster.a Block.o Matrix.o
ranlib libmaster.a

Building application for workers...
g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc -o Block.o
g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Matrix.cc -o Matrix.o
ar rvs libworker.a Block.o Matrix.o
ranlib libworker.a

...

Command successful. # The master for x86_64-linux-gnu compiled successfully

```

```

...

[ INFO ] Java libraries are searched in the directory: /usr/lib/jvm/java-1.8.0-openjdk-armhf/jre/lib/arm/
server
[ INFO ] Boost libraries are searched in the directory: /opt/install-arm/libboost

...

Building application for master...
arm-linux-gnueabi-g++ -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc Matrix.cc
ar rvs libmaster.a Block.o Matrix.o
ranlib libmaster.a

Building application for workers...
arm-linux-gnueabi-g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc -o
Block.o
arm-linux-gnueabi-g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Matrix.cc -
o Matrix.o
ar rvs libworker.a Block.o Matrix.o
ranlib libworker.a

...

Command successful. # The worker for arm-linux-gnueabi compiled successfully

...

[ INFO ] Java libraries are searched in the directory: /usr/lib/jvm/java-1.8.0-openjdk-amd64/jre/lib/amd64/
server
[ INFO ] Boost libraries are searched in the directory: /usr/lib/

...

Building application for master...
g++ -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc Matrix.cc
ar rvs libmaster.a Block.o Matrix.o
ranlib libmaster.a

Building application for workers...
g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Block.cc -o Block.o
g++ -DCOMPSS_WORKER -g -O3 -I. -I/Bindings/c/share/c_build/worker/files/ -c Matrix.cc -o Matrix.o
ar rvs libworker.a Block.o Matrix.o
ranlib libworker.a

...

Command successful. # The worker for x86_64-linux-gnu compiled successfully

```

*[The previous outputs has been cut for simplicity]*

Building for single architectures would lead to a directory structure quite different than the one obtained using the script for multiple architectures. In the single architecture case, only one master and one worker directories are expected. In the multiple architectures case, one master and one worker is expected per architecture.

```

.
|-- arm-linux-gnueabi
|   |-- worker
|   |   |-- gsbuild
|   |   |-- autom4te.cache
|-- src
|-- x86_64-linux-gnu
|   |-- master
|   |   |-- gsbuild

```



```

| |      '-- autom4te.cache
| '-- worker
|      '-- gsbuid
|      '-- autom4te.cache
'-- xml

```

(Note than only directories are shown).

### 4.3 Application Execution

The following environment variables must be defined before executing a COMPSs C/C++ application:

JAVA\_HOME: Java JDK installation directory (e.g. /usr/lib/jvm/java-8-openjdk/)

After compiling the application, two directories, master and worker, are generated. The master directory contains a binary called as the main file, which is the master application, in our example is called Matmul. The worker directory contains another binary called as the main file followed by the suffix “-worker”, which is the worker application, in our example is called Matmul-worker.

The *runcompss* script has to be used to run the application:

```

compss@bsc:~$ runcompss \
    --lang=c \
    -g \
    /home/compss/tutorial_apps/c/matmul_objects/master/Matmul 3 4 2.0

```

The complete list of options of the *runcompss* command is available in the *COMPSs User Manual: Application Execution* at <http://compss.bsc.es>.

### 4.4 Task Dependency Graph

Figure 1 depicts the task dependency graph for the Matmul application in its object version with 3x3 blocks matrices, each one containing a 4x4 matrix of doubles. Each block in the result matrix accumulates three block multiplications, i.e. three multiplications of 4x4 matrices of doubles.

The light blue circle corresponds to the initialization of matrix “A” by means of a method-task and it has an implicit synchronization inside. The dark blue circles correspond to the other two initializations by means of function-tasks; in this case the synchronizations are explicit and must be provided by the developer after the task call. Both implicit and explicit synchronizations are represented as red circles.

Each green circle is a partial matrix multiplication of a set of 3. One block from matrix “A” and the correspondent one from matrix “B”. The result is written in the right block in “C” that accumulates the partial block multiplications. Each multiplication set has an explicit synchronization. All green tasks are method-tasks and they are executed in parallel.

N = 3, Matrix size  
M = 4, Block size




Figure 1: Matmul Execution Graph.

## 5 Constraints

This section provides a detailed information about all the supported constraints by the COMPSs runtime for **Java**, **Python** and **C/C++** languages. The constraints are defined as key-value pairs, where the key is the name of the constraint. Table 8 details the available constraints names for *Java*, *Python* and *C/C++*, its value type, its default value and a brief description.

Java	Python	C / C++	Value type	Default value	Description
computingUnits	computing_units	ComputingUnits	<string>	"1"	Required number of computing units
processorName	processor_name	ProcessorName	<string>	"[unassigned]"	Required processor name
processorSpeed	processor_speed	ProcessorSpeed	<string>	"[unassigned]"	Required processor speed
processorArchitecture	processor_architecture	ProcessorArchitecture	<string>	"[unassigned]"	Required processor architecture
processorType	processor_type	ProcessorType	<string>	"[unassigned]"	Required processor type
processorPropertyName	processor_property_name	ProcessorPropertyName	<string>	"[unassigned]"	Required processor property
processorPropertyValue	processor_property_value	ProcessorPropertyValue	<string>	"[unassigned]"	Required processor property value
processorInternalMemorySize	processor_internal_memory_size	ProcessorInternalMemorySize	<string>	"[unassigned]"	Required internal device memory
processors	processors	-	List<@Processor>	"{"	Required processors (check Table 9 for Processor details)
memorySize	memory_size	MemorySize	<string>	"[unassigned]"	Required memory size in GBs
memoryType	memory_type	MemoryType	<string>	"[unassigned]"	Required memory type (SRAM, DRAM, etc.)
storageSize	storage_size	StorageSize	<string>	"[unassigned]"	Required storage size in GBs
storageType	storage_type	StorageType	<string>	"[unassigned]"	Required storage type (HDD, SSD, etc.)
operatingSystemType	operating_system_type	OperatingSystemType	<string>	"[unassigned]"	Required operating system type (Windows, MacOS, Linux, etc.)
operatingSystemDistribution	operating_system_distribution	OperatingSystemDistribution	<string>	"[unassigned]"	Required operating system distribution (XP, Sierra, openSUSE, etc.)
operatingSystemVersion	operating_system_version	OperatingSystemVersion	<string>	"[unassigned]"	Required operating system version
wallClockLimit	wall_clock_limit	WallClockLimit	<string>	"[unassigned]"	Maximum wall clock time
hostQueues	host_queues	HostQueues	<string>	"[unassigned]"	Required queues
appSoftware	app_software	AppSoftware	<string>	"[unassigned]"	Required applications that must be available within the remote node for the task

Table 8: Arguments of the *@constraint* decorator

All constraints are defined with a simple value except the *HostQueue* and *AppSoftware* constraints, which allow multiple values.

The *processors* constraint allows the users to define multiple processors for a task execution. This constraint is specified as a list of `@Processor` annotations that must be defined as shown in table 9

Annotation	Value type	Default value	Description
processorType	<string>	"CPU"	Required processor type (e.g. CPU or GPU)
computingUnits	<string>	"1"	Required number of computing units
name	<string>	"[unassigned]"	Required processor name
speed	<string>	"[unassigned]"	Required processor speed
architecture	<string>	"[unassigned]"	Required processor architecture
propertyName	<string>	"[unassigned]"	Required processor property
propertyValue	<string>	"[unassigned]"	Required processor property value
internalMemorySize	<string>	"[unassigned]"	Required internal device memory

Table 9: Arguments of the `@Processor` decorator

## 6 Known Limitations

The current COMPSs version (2.5) has the following limitations:

- **Exceptions:**

The current COMPSs version is not able to propagate exceptions raised from a task to the master. However, the runtime catches any exception and sets the task as failed.

- **Java tasks:**

Java tasks **must** be declared as **public**. Despite the fact that tasks can be defined in the main class or in other ones, we recommend to define the tasks in a separated class from the main method to force its public declaration.

- **Java objects:**

Objects used by tasks must follow the *java beans* model (implementing an empty constructor and getters and setters for each attribute) or implement the *serializable* interface. This is due to the fact that objects will be transferred to remote machines to execute the tasks.

- **Java object aliasing:**

If a task has an object parameter and returns an object, the returned value must be a new object (or a cloned one) to prevent any aliasing with the task parameters.

```
// @Method(declaringClass = "...")
// DummyObject incorrectTask (
//     @Parameter(type = Type.OBJECT, direction = Direction.IN) DummyObject a,
//     @Parameter(type = Type.OBJECT, direction = Direction.IN) DummyObject b
// );
public DummyObject incorrectTask (DummyObject a, DummyObject b) {
    if (a.getValue() > b.getValue()) {
        return a;
    }
    return b;
}

// @Method(declaringClass = "...")
// DummyObject correctTask (
//     @Parameter(type = Type.OBJECT, direction = Direction.IN) DummyObject a,
//     @Parameter(type = Type.OBJECT, direction = Direction.IN) DummyObject b
// );
public DummyObject correctTask (DummyObject a, DummyObject b) {
    if (a.getValue() > b.getValue()) {
        return a.clone();
    }
    return b.clone();
}

public static void main() {
    DummyObject a1 = new DummyObject();
    DummyObject b1 = new DummyObject();
    DummyObject c1 = new DummyObject();
    c1 = incorrectTask(a1, b1);
    System.out.println("Initial value: " + c1.getValue());
    a1.modify();
    b1.modify();
    System.out.println("Aliased value: " + c1.getValue());
}
```

```

DummyObject a2 = new DummyObject();
DummyObject b2 = new DummyObject();
DummyObject c2 = new DummyObject();
c2 = incorrectTask(a2, b2);
System.out.println("Initial value: " + c2.getValue());
a2.modify();
b2.modify();
System.out.println("Non-aliased value: " + c2.getValue());
}

```

- **Services types:**

The current COMPSs version only supports SOAP based services that implement the WS interoperability standard. REST services are not supported.

- **Use of file paths:**

The persistent workers implementation has a unique *Working Directory* per worker. That means that tasks should not use hardcoded file names to avoid file collisions and tasks misbehaviours. We recommend to use files declared as task parameters, or to manually create a sandbox inside each task execution and/or to generate temporary random file names.

- **Python constraints in the cloud:**

When using python applications with constraints in the cloud the minimum number of VMs must be set to 0 because the initial VM creation doesn't respect the tasks constraints. Notice that if no constraints are defined the initial VMs are still usable.

- **Intermediate files:**

Some applications may generate intermediate files that are only used among tasks and are never needed inside the master's code. However, COMPSs will transfer back these files to the master node at the end of the execution. Currently, the only way to avoid transferring these intermediate files is to manually erase them at the end of the master's code. Users must take into account that this only applies for files declared as task parameters and **not** for files created and/or erased inside a task.

- **User defined classes in Python:**

User defined classes in Python **must not be declared in the same file that contains the main method** (if `__name__ == '__main__'`) to avoid serialization problems of the objects.

- **Python object hierarchy dependency detection:**

Dependencies are detected only on the objects that are task parameters or outputs. Consider the following code:

```

# a.py
class A:
    def __init__(self, b):
        self.b = b
# main.py
from a import A
from pycompss.api.task import task
from pycompss.api.parameter import *

```

```

@task(obj = IN, returns = int)
def get_b(obj):
    return obj.b

@task(obj = INOUT)
def inc(obj):
    obj += [1]

def main():
    from pycomps.api import compss_wait_on
    my_a = A([5])
    inc(my_a.b)
    obj = get_b(my_a)
    obj = compss_wait_on(obj)
    print obj

if __name__ == '__main__':
    main()

```

Note that there should exist a dependency between `A` and `A.b`. However, PyCOMPSs is not capable to detect dependencies of that kind. These dependencies must be handled (and avoided) manually.

- **Python modules with global states:**

Some modules (for example `logging`) have internal variables apart from functions. These modules are not guaranteed to work in PyCOMPSs due to the fact that master and worker code are executed in different interpreters. For instance, if a `logging` configuration is set on some worker, it will not be visible from the master interpreter instance.

- **Python global variables:**

This issue is very similar to the previous one. PyCOMPSs does not guarantee that applications that create or modify global variables while worker code is executed will work. In particular, this issue (and the previous one) is due to Python's Global Interpreter Lock (GIL).

- **Python application directory as a module:**

If the Python application root folder is a python module (i.e: it contains an `__init__.py` file) then `runcompss` must be called from the parent folder. For example, if the Python application is in a folder with an `__init__.py` file named `my_folder` then PyCOMPSs will resolve all functions, classes and variables as `my_folder.object_name` instead of `object_name`. For example, consider the following file tree:

```

my_apps/
|- kmeans/
    |- __init__.py
    |- kmeans.py

```

Then the correct command to call this app is `runcompss kmeans/kmeans.py` from the `my_apps` directory.

- **Python early program exit:**

All intentional, premature exit operations must be done with `sys.exit`. PyCOMPSs



needs to perform some cleanup tasks before exiting and, if an early exit is performed with `sys.exit`, the event will be captured, allowing PyCOMPSs to perform these tasks. If the exit operation is done in a different way then there is no guarantee that the application will end properly.

- **Python with numpy and MKL:**

Tasks that invoke numpy and MKL may experience issues if tasks use a different number of MKL threads. This is due to the fact that MKL reuses threads along different calls and it does not change the number of threads from one call to another.

Please find more details on the COMPSs framework at

<http://compss.bsc.es>