Exascale Quantification of Uncertainties for
Technology and Science Simulation

D4.1 Working Python API to schedule MLC routines

Document information table

Contract number:	800898
Project acronym:	ExaQUte
Project Coordinator:	CIMNE
Document Responsible Partner:	CIMNE
Deliverable Type:	Software
Dissemination Level:	Public
Related WP & Task:	WP 4, Task 4.1
Status:	Final version

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No **800898**

Authoring

Prepared by:				
Authors	Partner	Modified Page/Sections	Version	Comments
Ramon Amela	BSC		V0	Creation of the document
Rosa M. Badia	BSC	Sections 1, 2, 3 and 5	V0	Changes in some sections
Riccardo Rossi	CIMNE			
Stanislav Bohm	IT4i			
Jakub Beranek	IT4i			
Contributors				

Change Log

Versions	Modified Page/Sections	Comments

Approval

Approved by:				
	Name	Partner	Date	OK
Task leader	Rosa M. Badia	BSC	27.7.18	OK
WP leader	Rosa M. Badia	BSC	27.7.18	OK
Coordinator	Riccardo Rossi	CIMNE	27.7.18	OK

Register for free at <https://www.scipedia.com> to download the version without the watermark

Executive summary

This deliverable focuses on the definition of a common API for the PyCOMPSs programming model and HyperLoom scheduler provided respectively by BSC and IT4I. The objective of the work is to hide the details of the actual task scheduling technology, so that the Multi Level Monte Carlo Python engine is agnostic of the backend being employed.

It includes the description of:

- Common API for calls
- Examples of usage
- Basic Documentation

The document also contains an initial description on how MPI-distributed data shall be treated from the scheduling point of view.

Register for free at <https://www.scipedia.com> to download the version without the watermark

Table of contents

1	Introduction	8
2	API definition	8
2.1	Common API	8
2.1.1	Launch process	8
2.1.2	Distribution strategy	8
2.1.3	API format	8
2.1.4	API calls	9
2.1.5	Task constraints	9
2.2	PyCOMPSs functionalities	9
2.2.1	Files as a special type	9
2.2.2	Constraint decorator	9
2.2.3	Implement decorator	12
2.3	HyperLoom functionalities	12
2.3.1	Dynamic resource requests	12
2.3.2	Run external program with zero-copy execution	13
3	Methods	13
3.1	Decorators	13
3.1.1	Task decorator	13
3.1.2	Constraint decorator	14
3.1.3	Implements	14
3.2	API calls	14
4	Example of usage	14
5	Data requirements	15
5.1	Single serialized object	15
5.2	Distributed serialized object	15
A	Methods	17
B	Example of usage	17

Register for free at <https://www.scipedia.com> to download the version without the watermark

List of Figures

1	Constraint basic usage	10
2	Constraint advanced usage	10
3	Implement usage example	13
4	Decorator definition for the local case	17
5	API calls definition for the local case	17
6	Basic use example	18
7	Monte Carlo step definition	19
8	Model definition	19
9	Function definition	20

Register for free at <https://www.scipedia.com> to download the version without the watermark

List of Tables

1	Arguments of the <i>@constraint</i> decorator	11
2	Arguments of the <i>@Processor</i> decorator	12

Register for free at <https://www.scipedia.com> to download the version without the watermark

Nomenclature / Acronym list

Acronym	Meaning
API	Application Programming Interface
ExaQute	EXAscale Quantification of Uncertainties for Technology and Science Simulation
DAG	Directed Acyclic Graph
FILE_IN	Path to a file passed to a function that is not modified
FILE_INOUT	Path to a file passed to a function that is modified during the call
FILE_OUT	Path to a file passed to a function that is created during the call
HPC	High Performance Computing
IN	Parameter of a function that is not modified
INOUT	Parameter of a function that is modified during the call
OpenMP	Open Multi Processing
MPI	Message Passing Interface
PBS	Portable Batch System
PyCOMPSs	Python binding for COMPS Superscalar
SLURM	Simple Linux Utility for Resource Management

The logo for Scipedia, featuring a stylized yellow 'S' and 'C' followed by the word 'IPEDIA' in grey. A large, semi-transparent watermark of this logo is overlaid on the page.

SCIPEDIA

Register for free at <https://www.scipedia.com> to download the version without the watermark

1 Introduction

PyCOMPSs and HyperLoom are task-based programming environment which enable the parallelization of sequential codes that can be then executed in distributed computing platforms.

State of the art description including references to relevant articles, e.g. [? ? ? ? ?]

2 API definition

The API has been designed in such a way that it is able to express the union of functionalities supported by both PyCOMPSs and HyperLoom. The aim of this section is to explain the API desirable properties and show the reasoning behind the definitions.

2.1 Common API

First of all, both frameworks has been analyzed to find the common points. All along the discussion, it has been taken into account that the users intended to use the API are not computer scientists. Hence, simplicity has been taken as a capital property to be achieved. In addition, the design tries to make as easy as possible the transformation of a sequential code into a distributed one.

2.1.1 Launch process

Both BSC and IT4I have already developed scripts that ease the deployment of the applications in HPC environments with queue systems. Hence, the user should be able to launch the application without managing all the information relative to the infrastructure. That is, being able to execute the code into different machines with a different amount of resources and a different architecture without changing it.

2.1.2 Distribution strategy

There are several distribution strategies. Nevertheless, both programming models define a workflow based on the *taskification* of the work. This means that some regions of code are defined to be executed remotely. The scheduler automatically detects the dependencies between tasks building the DAG that define the order in which the different tasks must be executed. In addition, generated objects are only brought to the master under demand.

2.1.3 API format

Considered the previous point, it has been decided to define the distributed portions of code with decorators. This strategy has several advantages. First of all, allows the user to define the main workflow sequentially, using a local definition of the decorator that avoids distributing the work. In addition, hides all the initialization process and avoids the direct communication between the user and the scheduler. Finally, makes possible to change from a sequential version to a distributed one o between both schedulers just changing the *module import*. This fact guarantees that the user can change the execution mode without modifying a single line of the code.

2.1.4 API calls

The most important functionalities shared between both programming models are the following ones:

- Task definition
- Get value to the master/client node
- Wait until all the tasks have finished their execution
- Delete object from the remote nodes

2.1.5 Task constraints

Both PyCOMPSs and HyperLoom allow the user to specify resource constraints. PyCOMPSs has a wider variety of resource characteristics that can be indicated by the user. On the other side, HyperLoom allows the user to define them inline a dynamic way.

2.2 PyCOMPSs functionalities

There are three main functionalities that has been added to the project since they were already supported by PyCOMPSs.

2.2.1 Files as a special type

A string passed to a function can be declared as a file. This fact implies that the Runtime copies the corresponding file to/between/from the worker nodes to ensure that a given task will find the file in the returned path that the Runtime pass to it as parameter. The same behavior than in the *object* case is offered, that is FILE_IN and FILE_INOUT. In addition, it is possible to indicate FILE_OUT in case the file is generated into the task. This behavior does not make sense in the *object* case since either a reference to an already created object is passed to the function or a new instance is created and given back to the main program as return parameter. Finally, it has to be considered that this separate consideration implies the creation of the special call `delete_file` to delete all the available copies in the worker nodes.

2.2.2 Constraint decorator

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 as shown in Figure 1.

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. Figure 2 shows how to express this constraints.

A full description of the supported constraints can be found in Table 1.

```
1 from pycompss.api.task import task
2 from pycompss.api.constraint import constraint
3 from pycompss.api.parameter import INOUT
4
5 %*{\bf @constraint }*)(ComputingUnits="4")
6 %*{\bf @task }*)(c = INOUT)
7 def func(a, b, c):
8     c += a*b
9     ...
```

Figure 1: Constraint basic usage

Register for free at <https://www.scipedia.com> to download the version without the watermark

```
1 from pycompss.api.task import task
2 from pycompss.api.constraint import constraint
3 from pycompss.api.parameter import INOUT
4
5 %*{\bf @constraint }*)(ComputingUnits="4", AppSoftware="numpy,scipy,gnuplot",
6     memorySize="$MIN_MEM_REQ")
7 %*{\bf @task }*)(c = INOUT)
8 def func(a, b, c):
9     c += a*b
10    ...
```

Figure 2: Constraint advanced usage

Field	Value type	Default value	Description
ComputingUnits	<string>	"1"	Required number of computing units
ProcessorName	<string>	"[unassigned]"	Required processor name
ProcessorSpeed	<string>	"[unassigned]"	Required processor speed
ProcessorArchitecture	<string>	"[unassigned]"	Required processor architecture
ProcessorType	<string>	"[unassigned]"	Required processor type
ProcessorPropertyName	<string>	"[unassigned]"	Required processor property
ProcessorPropertyValue	<string>	"[unassigned]"	Required processor property value
ProcessorInternalMemorySize	<string>	"[unassigned]"	Required internal device memory
-	List<@Processor>	"{}"	Required processors (check Table 2 for Processor details)
MemorySize	<string>	"[unassigned]"	Required memory size in GBs
MemoryType	<string>	"[unassigned]"	Required memory type (SRAM, DRAM, etc.)
StorageSize	<string>	"[unassigned]"	Required storage size in GBs
StorageType	<string>	"[unassigned]"	Required storage type (HDD, SSD, etc.)
OperatingSystemType	<string>	"[unassigned]"	Required operating system type (Windows, MacOS, Linux, etc.)
OperatingSystemDistribution	<string>	"[unassigned]"	Required operating system distribution (XP, Sierra, openSUSE, etc.)
OperatingSystemVersion	<string>	"[unassigned]"	Required operating system version
WallClockLimit	<string>	"[unassigned]"	Maximum wall clock time
HostQueues	<string>	"[unassigned]"	Required queues
AppSoftware	<string>	"[unassigned]"	Required applications that must be available within the remote node for the task

Table 1: 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 2

Annotation	Value type	Default value	Description
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
type	<string>	"[unassigned]"	Required processor type
propertyName	<string>	"[unassigned]"	Required processor property
propertyValue	<string>	"[unassigned]"	Required processor property value
internalMemorySize	<string>	"[unassigned]"	Required internal device memory

Table 2: Arguments of the *@Processor* decorator

2.2.3 Implement decorator

PyCOMPSs allows the user to define several versions of the same tasks. The main idea behind this functionality is having several ways to perform the same operation. In the general case, the constraints of each version are different.

For example, it would be possible to have an implementation for an accelerator (GPU or FPGA) and an implementation using common CPUs. It would also be possible to define an implementation that uses **OpenMP** occupying a single node and a **MPI** version occupying several computing nodes.

Figure 3 a use example of this decorator. In this case, the secondary implementation uses a library that may be installed in a subset of the computing nodes. The information that must be provided is the class in which the main implementation is coded and the method name.

2.3 HyperLoom functionalities

2.3.1 Dynamic resource requests

HyperLoom allows to define resource request for each task individually.

```
1 from pycompss.api.implement import implement
2
3 @implement(source_class="sourcemodule", method="main_func")
4 @constraint(AppSoftware="numpy")
5 @task(returns=list)
6 def myfunctionWithNumpy(list1, list2):
7     # Operate with the lists using numpy
8     return resultList
9
10 @task(returns=list)
11 def main_func(list1, list2):
12     # Operate with the lists using built-int functions
13     return resultList
```

Figure 3: Implement usage example

2.3.2 Run external program with zero-copy execution

HyperLoom provides "run" method to running external programs. Data objects are managed via symlinking and moving on RAM disk. It avoids unnecessary copy of data and reading data by worker at all if computation stays on the same worker.

3 Methods

Both frameworks have to implement all the decorations and API calls described in this section even if they don't have any impact in the scheduling process. This fact guarantees that the code do not crash when switching between them. It has to be taken into account that Figures 4 and 5 contain the wrapper coded to work in local. This code helps understanding the expected behavior but does not have any scheduling implementation.

3.1 Decorators

3.1.1 Task decorator

First of all, a decorator has been defined to indicate the tasks. It's basic syntax is shown on Figure 4. All the functions marked with this decorator will be executed remotely. By default it is assumed that all the function parameters are IN and there is no return value. In addition, there are two hints that must be given otherwise:

1. **returns** allows the user to say how many return values has the function.
2. **variable_name={INOUT, FILE_IN, FILE_OUT, FILE_INOUT}** makes possible to indicate either that the input object is modified or that the input string corresponds to a file. This way, the scheduler can transfer the file between the different worker nodes to make the execution possible.

Finally, it has to be taken into account that there is an optional parameter associated to the keyword **scheduling_constraints** that makes possible to pass an instance of **ExaQuteExecutionConstraints** that allows the user to specify task constraints in a dynamic way. That is, in execution time. This point has to be discussed more deeply considering the users feedback. Even if it allows them more expresivity, introduces scheduling dependent code into the user's code, fact that was agreed to avoid at the maximum. Alternatively, it would be possible to define several versions of the same call with different execution constraints.

3.1.2 Constraint decorator

Considering that this is a functionality that BSC had already implemented, the syntax remains exactly as presented on subsection 2.2. More keywords could be added considering the user's demands and the project's particularities.

3.1.3 Implements

This case is analogous to the previous point. Hence, the same notation already implemented on PyCOMPSs has been kept.

3.2 API calls

Figure 5 shows the basic definition of the api calls implemented. The following points summary their basic expected behavior:

- `barrier()`
The purpose of this call is to wait until all the tasks have been executed.
- `get_value_from_remote(obj)`
This API call brings back to the master/client the value of a remote computed object.
- `delete_object(obj)`
This API call removes all the copies allocated in the workers of the given object.
- `delete_file(path_to_file)`
This API call removes all the copies allocated in the workers of the given file.
- `compute(value)`
This API call submits the given value to the scheduler, starting the computation considering all the DAG generated until the moment.

4 Example of usage

This example, even is a toy code that does not perform any useful computation, has the property of testing all the functionalities described in section 3. Figure 6 shows how the example code launches the initialization of several models with really variate granularity in distributed nodes. This generated model is defined in Figure 8. The most important think to realize that it contains the **Kratos** model that will have to be handled in the real project. In addition, an intermediate class has been defined as shown in Figure 7 to show how distributed computations can be hidden in an intermediate class to ease the workflow definition. All the distributed functions have been defined in Figure 7 to ease the identification of the parts that will be executed in the worker nodes. The interactions with the scheduler are completely contained in this file. Hence, changing from local to distributed and between both frameworks is done changing a single import.

5 Data requirements

Concerning the data requirements, it is possible to define two different project phases. Initially, data will be serialized to a single string. Next, distributed objects will be handled.

5.1 Single serialized object

In the first stage, medium size tasks will be launched in such a way that all the data can be fit in the memory of a single node. All the data passed to a task as a parameter must be serializable. This fact implies the definition of the following functions:

1. *save*

It is possible to define this functions in two different ways.

First of all, the function could have the following parameters:

- *obj*: input object to serialize
- *return*: a string or a stream in which to write the serialized object

Otherwise, an alternative definition is also possible:

- *obj*: input object to serialize
- *path*: input string representing the path into which a file containing the serialized object should be created

2. *load*

Afterwards, the serialized objects must be recoverable. This implies providing the symmetrical functionalities to the ones defined previously.

In case the first option has been chosen, the following parameters are involved into the function call:

- *str*: input string or stream containing the serialized object
- *return*: original object passed to the *save* function

The second version of this function should have the following parameters:

- *path*: input string representing the path of the file containing the serialized object
- *return*: original object passed to the *save* function

5.2 Distributed serialized object

In a more advanced stage, *mpi* simulations will be run, occupying several computing nodes. The case in which a single node memory is not capable to store all the data handled by each task is contemplated. Hence, a single serialized string/file is no longer possible. There is a deliverable in the 18th month concerning this problematic. Nevertheless, the discussion has already started. Some desirable properties have already been defined:

- The objects should be serialized in a distributed way, meaning that each *MPI* process should be responsible to the serialization and later recovery of its piece of data

- Since there is no longer a single serialization for each object, two main strategies are possible to express the dependencies between tasks:
 1. Define a component into the programming model that can be directly called from the mpi processes in such a way that it tracks all the serialized model parts and its location, being in charge to store it in a file in the early stages and in memory in more advanced phases
 2. Not considering the whole model at all and define all the dependencies in function of the several parts. In this case, the scheduler should be aware of the *MPI* process associated to each dependency in order to perform the transferences and set all the *MPI* environment variables accordingly.

Both solutions succeed in avoiding the shared file system and the serialization of the whole model into a single serialization. However, the first option seems more interesting as it delegates all the data handling part to the programming models. It keeps better the philosophy of having as less code as possible related to its distribution.

Nevertheless, it is important to keep in mind that this phase is still in a really early stage. Some discussion between all the involved partners needs to be carried out. A consensus should be achieved in such a way that the desired properties are achieved and the solution fits as well as possible in the already implemented codes.


```
1 class Exaquite_task(object):
2
3     def __init__(self, *args, **kwargs):
4         pass
5
6     def __call__(self, f):
7         def g(*args, **kwargs):
8             if "scheduling_constraints" in kwargs:
9                 del kwargs["scheduling_constraints"]
10            return f(*args, **kwargs)
11         return g
12
```

Figure 4: Decorator definition for the local case

```
1 def barrier():
2     pass
3
4 def get_value_from_remote(obj):
5     return obj
6
7 def delete_object(obj):
8     del obj
9
10 def delete_file(file_path):
11     import os
12     os.remove(file_path)
13
14 def compute(obj):
15     return obj
16
```

Figure 5: API calls definition for the local case

References

A Methods

B Example of usage

```

1 import sys
2 import MultilevelMonteCarloFunctions
3 from MultilevelMonteCarloStep import MultilevelMonteCarloStep
4
5 short = False
6
7 if short:
8     multiplier_low = 1
9     multiplier_up = 2
10 else:
11     multiplier_low = 13
12     multiplier_up = 20
13
14 if __name__ == "__main__":
15     models_path = sys.argv[1]
16
17     steps = MultilevelMonteCarloFunctions
18         .generate_models(models_path, 3, [], short, multiplier_low, multiplier_up)
19
20     while len(steps) > 0:
21         runs = []
22         comparisons = []
23         versions = []
24         for i in xrange(len(steps)):
25             runs.append([])
26             if i > 0:
27                 comparisons.append([])
28                 for j in xrange(steps[i][1].amount_executions):
29                     current_step_model = MultilevelMonteCarloStep(steps[i][0] + ".mdpa")
30
31                     steps[i-1][1].cpus_per_task = 2
32                     versions.append(current_step_model.generate_random(scheduling_
33 constraints = steps[i-1]))
34
35                     current_step_model.run()
36
37                     runs[i].append(current_step_model)
38                     if i > 0:
39                         comparisons[i-1].append(MultilevelMonteCarloFunctions.compute(runs[i
40 ][j].compare_models(runs[i-1][j])))
41
42                     if(len(steps) > 1):
43                         partial_list = []
44                         for compare_list in comparisons:
45                             partial_list.append(MultilevelMonteCarloFunctions.reduce_list(compare_
46 list))
47
48                     verification_value = MultilevelMonteCarloFunctions.reduce_list(partial_list
49 )
50
51                     verification_value = MultilevelMonteCarloFunctions.get_value(verification_
52 value)
53
54                     versions_count = MultilevelMonteCarloFunctions.count_versions(*versions)
55                     versions_count = MultilevelMonteCarloFunctions.get_value(versions_count)
56
57                     print("VERIFICATION VALUE: " + str(verification_value.params))
58                     print("AMOUNT OF DIFFERENT VERSIONS USED: " + str(versions_count))
59
60                 else:
61                     verification_value = 0
62
63                 multiplier_low /= 2
64                 multiplier_up /= 2
65                 if short:
66                     steps = MultilevelMonteCarloFunctions
67                         .generate_models(models_path, 2, verification_value, short, multiplier_
68 low, multiplier_up)
69                 else:
70                     steps = MultilevelMonteCarloFunctions
71                         .generate_models(models_path, 3, verification_value, short, multiplier_
72 low, multiplier_up)

```

Figure 6: Basic use example

```
1 import MultilevelMonteCarloFunctions
2
3 class MultilevelMonteCarloStep(object):
4
5     def __init__(self, model):
6         self.model = MultilevelMonteCarloFunctions.load_model(model)
7         self.params = None
8
9     def generate_random(self, scheduling_constraints = None):
10         self.params = MultilevelMonteCarloFunctions.generate_random(self.model,
11                               scheduling_constraints = scheduling_constraints)
12         return self.params
13
14     def run(self):
15         MultilevelMonteCarloFunctions.run(self.model, self.params)
16
17     def compare_models(self, modelB):
18         return MultilevelMonteCarloFunctions.compare_models(self.model, modelB.model)
19
```

Figure 7: Monte Carlo step definition

```
1 import KratosMultiphysics
2 import KratosMultiphysics.FluidDynamicsApplication
3
4 class MultilevelMonteCarloModel(object):
5
6     def __init__(self, path_to_model = None):
7         if(not path_to_model is None):
8             model_part_name = "MainRestart"
9             model_part = KratosMultiphysics.ModelPart(model_part_name)
10             model_part.AddNodalSolutionStepVariable(KratosMultiphysics.DISPLACEMENT)
11             model_part.AddNodalSolutionStepVariable(KratosMultiphysics.VISCOSITY)
12             model_part_io = KratosMultiphysics.ModelPartIO(path_to_model)
13             model_part_io.ReadModelPart(model_part)
14             self.model_part = model_part
15
16     def set_model_part(self, model_part):
17         self.model_part = model_part
18
```

Figure 8: Model definition

```

1  import random
2  from MultilevelMonteCarloParams import MultilevelMonteCarloParams
3  from MultilevelMonteCarloModel import MultilevelMonteCarloModel
4  import KratosMultiphysics
5  import math
6  from subprocess import Popen, PIPE
7  from ExaquiteExecutionConstraints import ExaquiteExecutionConstraints
8
9  from ExaquiteTaskPyCOMPSs import *
10
11 def _listdir_shell(path, *lsargs):
12     p = Popen(('ls', path) + lsargs, shell=False, stdout=PIPE, close_fds=True)
13     return [path.rstrip('\n') for path in p.stdout.readlines()]
14
15 @ExaquiteTask(path_to_model = FILE_IN, returns = 1)
16 def load_model(path_to_model):
17     return MultilevelMonteCarloModel(path_to_model[:-5])
18
19 @constraint(computingUnits = "2")
20 @ExaquiteTask(model = INOUT, returns = 1)
21 def generate_random(model):
22     model_part = model.model_part
23     # generate random number
24     for node in model_part.Nodes:
25         rand = random.random()
26         node.SetSolutionStepValue(KratosMultiphysics.VISCOSITY, 0, rand)
27     params = MultilevelMonteCarloParams(1)
28     return params
29
30 @implement(source_class="MultilevelMonteCarloFunctions", method="run")
31 @ExaquiteTask(model = INOUT, params = INOUT)
32 def run_2(model, params):
33     model_part = model.model_part
34     for node in model_part.Nodes:
35         node.SetSolutionStepValue(KratosMultiphysics.VISCOSITY, 0,
36                                   node.GetSolutionStepValue(KratosMultiphysics.VISCOSITY, 0) + 1.0)
37     params.set_value(2)
38
39 @ExaquiteTask(model = INOUT, params = INOUT)
40 def run(model, params):
41     model_part = model.model_part
42     for node in model_part.Nodes:
43         node.SetSolutionStepValue(KratosMultiphysics.VISCOSITY, 0,
44                                   node.GetSolutionStepValue(KratosMultiphysics.VISCOSITY, 0) + 1.0)
45     params.set_value(1)
46
47 @ExaquiteTask(returns = 1)
48 def compare_models(model_1, model_2):
49     model_part_1 = model_1.model_part
50     model_part_2 = model_2.model_part
51     # take the average
52     avgModel = 0.0
53     n = max(1, len(model_part_1.Nodes) + len(model_part_2.Nodes))
54     for node in model_part_1.Nodes:
55         avgModel += node.GetSolutionStepValue(KratosMultiphysics.VISCOSITY)
56     for node in model_part_2.Nodes:
57         avgModel += node.GetSolutionStepValue(KratosMultiphysics.VISCOSITY)
58     avgModel /= n
59     return MultilevelMonteCarloParams(avgModel)
60
61 def get_value(obj):
62     return get_value_from_remote(obj)
63
64 @ExaquiteTask(returns = 1)
65 def reduce_function(*reduce_list):
66     sum_list = reduce_list[0].params
67     for elem in reduce_list[1:]:
68         sum_list += elem.params
69     avg = sum_list / len(reduce_list)
70     return MultilevelMonteCarloParams(avg)
71
72 @ExaquiteTask(returns = 1)
73 def count_versions(*versions):
74     from collections import Counter
75     int_list = []
76     for elem in versions:
77         int_list.append(elem.params)
78     ret = dict(Counter(int_list))
79     return ret
80

```

Figure 9: Function definition