

---

# **Decimate Documentation**

*Release 0.9.5*

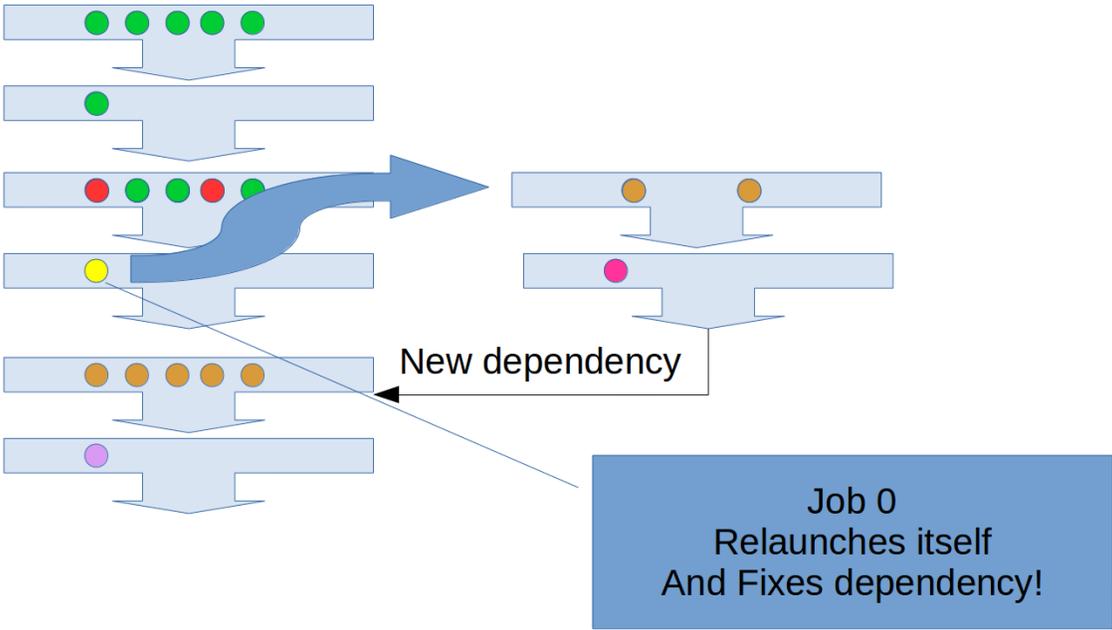
**Samuel KORTAS**

**Apr 10, 2018**



<b>1</b>	<b>What is <i>Decimate</i>?</b>	<b>3</b>
1.1	Features . . . . .	3
1.2	Automated restart in case of failure . . . . .	4
1.3	Fully user configurable environment . . . . .	4
<b>2</b>	<b>Installation</b>	<b>5</b>
2.1	Requirements . . . . .	5
2.2	Distribution . . . . .	5
2.3	Installing <i>Decimate</i> using PIP . . . . .	5
2.3.1	Installing <i>Decimate</i> as root using PIP . . . . .	5
2.3.2	Installing <i>Decimate</i> as user using PIP . . . . .	6
2.4	Installing <i>Decimate</i> using Anaconda . . . . .	6
2.5	Source . . . . .	6
<b>3</b>	<b>Using <i>Decimate</i></b>	<b>7</b>
3.1	Submitting a Workflow . . . . .	7
3.1.1	Submitting a job . . . . .	7
3.1.2	options . . . . .	7
3.1.3	single job . . . . .	7
3.1.4	job depending on previous submitted jobs . . . . .	8
3.1.5	submitting a job with a user-defined checking function . . . . .	8
3.1.6	other kinds of workflows . . . . .	8
3.2	Checking the current status . . . . .	9
3.3	Displaying the log file . . . . .	9
3.4	Cancelling the whole workflow . . . . .	9
<b>4</b>	<b>Examples of Workflows</b>	<b>11</b>
4.1	Test job . . . . .	11
4.2	Nominal 2 job workflow . . . . .	11
4.3	parametric job workflow . . . . .	12
<b>5</b>	<b>Fault-tolerant Workflows</b>	<b>13</b>
5.1	Adding a user-defined checking function . . . . .	13
5.2	Successful job submission . . . . .	14
5.3	Failed job submission and automated restarting . . . . .	14
5.4	Setting the number of restart . . . . .	16

<b>6</b>	<b>Parameters combination</b>	<b>17</b>
6.1	array of values . . . . .	17
6.2	Combined parameter sweep . . . . .	18
6.3	Parameters depending on simple formulas . . . . .	18
6.4	More complex Python expressions . . . . .	19
<b>7</b>	<b>Shell API</b>	<b>21</b>
7.1	dbatch . . . . .	21
7.2	dstat . . . . .	22
7.3	dlog . . . . .	22
7.4	dkill . . . . .	22
7.5	environment variables . . . . .	22
7.6	Job script directives . . . . .	22





---

## What is *Decimate*?

---

Developed by the KAUST Supercomputing Laboratory (KSL), *Decimate* is a SLURM extension written in Python allowing the user to handle jobs per hundreds in an efficient and transparent way. In this context, the constraint limiting the number of jobs per users is completely masked. The time consuming burden of managing thousands of jobs by hand is also alleviated by making available to the user the concept of workflow gathering a set of jobs that he can manipulate as a whole.

*Decimate* is released as an Open Source Software under BSD Licence. It is available at

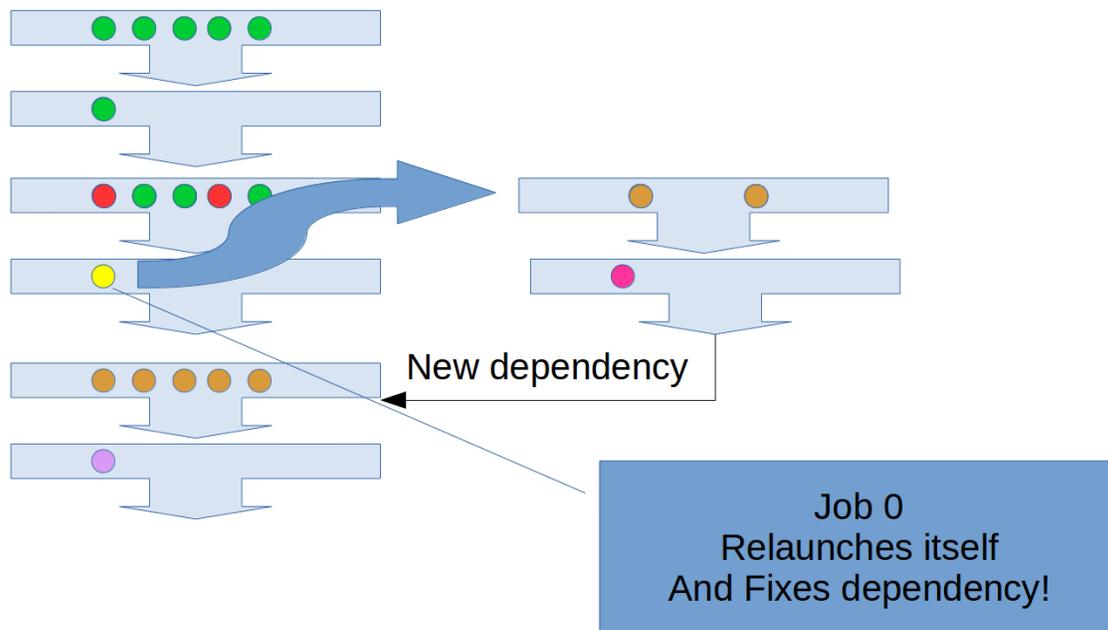
### 1.1 Features

*Decimate* allows a user to:

- Submit any number of jobs regardless of any limitation set in the scheduling policy on the maximum number of jobs authorized per user.
- Manage all the submitted jobs as a single workflow easing their submission, monitoring, deletion or reconfiguration.
- Ease the definition, submission and management of jobs run on a large set of combinations of parameters.
- Benefit from a centralized log file, a unique point of capture of relevant information about the behavior of the workflow. From Python or shell, at any time and from any jobs, the logging levels info, debug, console and mail are available.
- Send fully-configurable mail messages detailing the current completion of the workflow at any step of its execution.
- Easily define a procedure (in shell or Python) to check for correctness of the results obtained at the end of given step. Having access to the complete current status of the workflow, this procedure can make the decision on-the-fly either to stop the whole workflow, to resubmit partially the failing components as is, or to modify it dynamically.

## 1.2 Automated restart in case of failure

In case of failure of one part of the workflow, *Decimate* automatically detects the failure, signals it to the user and launches the misbehaving part after having fixed the job dependency. By default if the same failure happens three consecutive times, *Decimate* cancels the whole workflow removing all the depending jobs from the scheduling. In a next version, *Decimate* will allow the automatic restarting of the workflow once the problem causing its failure has been cured.



## 1.3 Fully user configurable environment

*Decimate* also allows the user to define his own mail alerts that can be sent at any point of the workflow.

Some customized checking functions can also be designed by the user. Their purpose is to validate if a step of the workflow was successful or not. It could involve checking for the presence of some result files, grepping some error or success messages in them, computing ratio or checksum. . . These intermediate results can be easily transmitted to *Decimate* validating or not the correctness of any step. They can also be forwarded by mail to the user where as the workflow is executing.

### 2.1 Requirements

*Decimate* should work with any cluster based on Unix operating systems which provides Python 2.7 and using SLURM as a scheduler. It also depends on the python packages numpy, pandas and clustershell.

In a further release, *Decimate* is planned to be compatible with Python 3 and no dependency on numpy will be imposed.

### 2.2 Distribution

*Decimate* is an open-source project distributed under the BSD 2-Clause “Simplified” License which means that many possibilities are offered to the end user including the fact to embed *Decimate* in one own software.

Its stable production branch is available via github at <https://github.com/KAUST-KSL/decimate>, but its latest production and development branch can be found at <https://github.com/samkos/decimate>

The most recent documentation about *Decimate* can be browsed at <http://decimate.readthedocs.io>.

### 2.3 Installing *Decimate* using PIP

#### 2.3.1 Installing *Decimate* as root using PIP

To install *Decimate* as a standard Python package using PIP<sup>1</sup> as root:

```
$ pip install decimate
```

Or alternatively, using the source tarball:

---

<sup>1</sup> pip is a tool for installing and managing Python packages, such as those found in the Python Package Index

```
$ pip install decimate-0.9.x.tar.gz
```

### 2.3.2 Installing *Decimate* as user using PIP

To install *Decimate* as a standard Python package using PIP as an user:

```
$ pip install --user decimate
```

Or alternatively, using the source tarball:

```
$ pip install --user decimate-0.9.x.tar.gz
```

Then, you just need to update your PYTHONPATH environment variable to be able to import the library and PATH to easily use the tools:

```
$ export PYTHONPATH=$PYTHONPATH:~/.local/lib  
$ export PATH=$PATH:~/.local/bin
```

Configuration files are installed in `~/.local/etc/decimate` and are automatically loaded before system-wide ones (for more info about supported user config files, please see the `decimate-config` config section).

## 2.4 Installing *Decimate* using Anaconda

*Decimate* is also available in Anaconda from the *hpc4all* channel. It can be installed with the command:

```
$ conda install -c hpc4all decimate
```

## 2.5 Source

Current source is available on Github, use the following command to retrieve the latest stable version from the repository:

```
$ git clone -b prod git@github.com:samkos/decimate.git
```

and for the development version:

```
$ git clone -b dev git@github.com:samkos/decimate.git
```

Via *Decimate*, four commands are added to the user environment: **dbatch** to submit workflows, **dstat** to monitor their current status, **dlog** to tail the log information produced and **dkill** to cancel the execution of the workflow.

## 3.1 Submitting a Workflow

For *Decimate*, a *workflow* is a set of jobs submitted from a same directory. These jobs can depend on one another and be job array of any size.

### 3.1.1 Submitting a job

### 3.1.2 options

*Decimate* **dbatch** command accepts the same *options* as the SLURM **sbatch** command and extends it in two ways:

- it transparently submits the user job within a fault-tolerant framework
- it adds new options to manage the workflow execution if a problem occurs
  - `--check=SCRIPT_FILE` points to a user script (either in python or shell) to validate the correctness of the job at the end of its execution
  - `--max-retry=MAX_RETRY` setting number of time a step can fail and be restarted automatically before failing the whole workflow (3 per default)

### 3.1.3 single job

Here is how to submit a simple job:

```
dbatch --job-name=job_1 my_job.sh
```

```
[MSG ] submitting job job_1 (for 1) --> Job # job_1-0-1 <-depends-on None
[INFO ] launch-0!0:submitting job job_1 [1] --> Job # job_1-0-1 <-depends-on None
Submitted batch job job_1-0-1
[1] --> Job # job_1-0-1 <-depends-on None
```

Notice how the command syntax is similar to **sbatch** command.

- In lines starting with [MSG], [INFO], or [DEBUG], *Decimate* gives us additional information about what is going on.
- All the traces [INFO], or [DEBUG] also appears in the corresponding job output file as well as in *Decimate* central log file dumped in `<current_directory>/decimate/LOGS/decimate.log` [MSG] traces only appears at the console or in the output file of the job.
- for *Decimate*, every job is considered as a job array. In this simple case, it considers an array of job made of a single element 1-1. In the traces, the array indice shows in “(for 1)”, “submitting job job\_1 [1]”, or “job job\_1-0-1”. (if needed check [SLURM job array documentation](#) for more information).
- Every job submitted via *Decimate* is part of a fault-tolerant environment. At the end of its execution, its correctness is systematically checked thanks to a user defined function or by default thanks the return code of the job given by SLURM. If the job is not considered as correct, (and if the return code of the user-defined function is not *ABORT*), the job is automatically resubmitted for a first and a second attempt if needed. In the traces, the attempt number shows as the second figure in the job denomination: “job job\_1-0-1”.

### 3.1.4 job depending on previous submitted jobs

Here is how to submit a job depending on a previous job:

```
dbatch --dependency=job_1 --job-name=job_2 my_job.sh
[INFO ] launch-0!0:Workflow has already run in this directory, trying to continue it
[MSG ] submitting job job_2 (for 1) --> Job # job_2-0-1 <-depends-on 218459
[INFO ] launch-0!0:submitting job job_2 [1] --> Job # job_2-0-1 <-depends-on 218459
Submitted batch job job_2-0-1
[1] --> Job # job_2-0-1 <-depends-on 218459
```

It again matches **sbatch** original syntax with the subtlety that via *Decimate* dependency can be expressed with respect to a previous job name and not only to a previous job id as **SLURM** only allows it.

- It makes it more convenient to write automated script.
- At this submission time, *Decimate* checks if a previous submitted job has actually been submitted with this particular name. If not, an error will be issued and the submission is canceled.
- Of course, dependency on a previous job id is also supported.

### 3.1.5 submitting a job with a user-defined checking function

Fault-tolerant jobs submission and behavior is addressed in [Fault-tolerant Workflows](#).

### 3.1.6 other kinds of workflows

A comprehensive list of job examples can be found in [Examples of Workflows](#).

## 3.2 Checking the current status

The current workflow status can be checked with **dstat**:

```
dstat
```

When no job has been submitted from the current directory. **dstat** shows:

```
[MSG ] No workflow has been submitted yet
```

When jobs submitted submitted the current directory are currently running . **dstat** shows:

```
[MSG ] step job_1-0:1-1          SUCCESS  SUCCESS: 100%  FAILURE:  0%  ->└
↔ []
[MSG ] step job_2-0:1-1          RUNNING   SUCCESS:   0%  FAILURE:  0%  ->└
↔ []
```

And when a workflow is completed:

```
dstat
[MSG ] CHECKING step : job_2-0 task 1
[MSG ] step job_1-0:1-1          SUCCESS  SUCCESS: 100%  FAILURE:  0%  ->└
↔ []
[MSG ] step job_2-0:1-1          SUCCESS  SUCCESS: 100%  FAILURE:  0%  ->└
↔ []
```

## 3.3 Displaying the log file

The current *Decimate* log file can be checked with **dlog**:

```
dlog
```

## 3.4 Cancelling the whole workflow

The current workflow can be completely killed with the command **dkill**:

```
dkill
```

If no job of the workflow is either running, queueing or waiting to be queued, **dkill** prints:

```
[INFO ] No jobs are currently running or waiting... Nothing to kill then!
```

If any job is still waiting or running, *dkill* asks a confirmation to the user and cancels all jobs from the current workflow.



### 4.1 Test job

Let *my\_job.sh* be the following example job:

```
#!/bin/bash
#SBATCH -n 1
#SBATCH -t 0:05:00

echo job running on...
hostname
sleep 10

echo job DONE
```

If not done yet, we first load the *Decimate* module:

```
module load decimate
```

### 4.2 Nominal 2 job workflow

Then submission of jobs follows the same syntax than with the **sbatch** command:

```
dbatch --job-name=job_1 my_job.sh
```

```
[MSG ] submitting job job_1 (for 1) --> Job # job_1-0-1 <-depends-on None
[INFO ] launch-0!0:submitting job job_1 [1] --> Job # job_1-0-1 <-depends-on None
Submitted batch job job_1-0-1
[1] --> Job # job_1-0-1 <-depends-on None
```

```

dbatch --dependency=job_1 --job-name=job_2 my_job.sh
[INFO ] launch-0!0:Workflow has already run in this directory, trying to continue it
[MSG ] submitting job job_2 (for 1) --> Job # job_2-0-1 <-depends-on 218459
[INFO ] launch-0!0:submitting job job_2 [1] --> Job # job_2-0-1 <-depends-on 218459
Submitted batch job job_2-0-1
[1] --> Job # job_2-0-1 <-depends-on 218459

```

```
dstat
```

```

[MSG ] step job_1-0:1-1                SUCCESS  SUCCESS: 100%  FAILURE:  0%  ->  
↪ []
[MSG ] step job_2-0:1-1                RUNNING   SUCCESS:   0%  FAILURE:  0%  ->  
↪ []

```

```

dstat
[MSG ] CHECKING step : job_2-0 task 1
[INFO ] launch-0!0:no active job in the queue, changing all WAITING in ABORTED???
[MSG ] step job_1-0:1-1                SUCCESS  SUCCESS: 100%  FAILURE:  0%  ->  
↪ []
[MSG ] step job_2-0:1-1                SUCCESS  SUCCESS: 100%  FAILURE:  0%  ->  
↪ []

```

### 4.3 parametric job workflow

Then submission of parametric jobs follows the same syntax than with the **sbatch** command adding a reference to a text file describing the set of parameters to be tested:

```
dbatch --job-name=job_1 -P parameters.txt my_job.sh
```

How to build the file *parameters.txt* is described at [Parameters combination](#).

## 5.1 Adding a user-defined checking function

*Decimate* allows the user to define its own function to qualify a job as *ABORT*, *SUCCESS* or *FAILED*. This can be a simple bash script file or a program written in Python. For example here is a typical script *check\_job.sh* written in shell checking if the message 'job DONE' appears in the job output file:

```
job_step=$1
attempt=$2
task_id=$3
running_dir=$4
output_file=$5
error_file=$5
is_job_completed=$6

echo job_step=$job_step attempt=$attempt task_id=$task_id
echo running_dir=$running_dir
echo output_file=$output_file
echo error_file=$error_file
echo is_job_completed=$is_job_completed

SUCCESS=0
FAILURE=-1
ABORT=-9999

grep 'job DONE' $output_file
```

All the parameters are passed to the script as arguments added on the command line.

## 5.2 Successful job submission

When submitting the job, one only adds `-check` followed by the path of the checking job script

```
dbatch --check=check_job.sh --job-name=job_1 my_job.sh
```

`my_job.sh` is the following job which will be checked as successful because echoing the string `job DONE`:

```
#!/bin/bash
#SBATCH -p debug
#SBATCH -n 1
#SBATCH -t 0:01:00

echo job running on...
hostname
sleep 10

echo job DONE
```

One can follow the current status of the workflow thanks to `dlog` that displays the log file attached to the current workflow.

```
dlog
()
=====
Currently Tailing ...
/home/kortass/DECIMATE-GITHUB/.decimate/LOGS/decimate.log
Hit CTRL-C to exit... Hit CTRL-C to exit...
=====
...
[INFO ] launch-0!0:submitting job 1 (for 1) --> Job # 1-0-1 <-depends-on None
[INFO ] launch-0!0:submitting job chk_1 (for 1) --> Job # chk_1-0-1 <-depends-on 1-0-
↪1
[INFO ] chk_1-1!0: ok everything went fine for the step 1 (1) --> Step chk_1 (1) is_
↪starting... @ (2018-02-21 11:31:06)
[INFO ] chk_1-1!0:===== workflow is finishing ===== @ (2018-02-21_
↪11:31:09)
```

## 5.3 Failed job submission and automated restarting

In the case of failure, here is what is observed when submitting a job that fails:

```
dbatch --check=check_job.sh --job-name=job_1 my_job_failed.sh
```

`my_job_failed.sh` is the following job which will be checked as failed because echoing not the string `job DONE`:

```
#!/bin/bash
#SBATCH -p debug
#SBATCH -n 1
#SBATCH -t 0:01:00

echo job running on...
hostname
```

```
echo job FAILED
```

which leads to the following results observed with the command **dlog** that displays the log file attached to the current workflow.

```
dlog
()
=====
Currently Tailing ...
/home/kortass/DECIMATE-GITHUB/.decimate/LOGS/decimate.log
Hit CTRL-C to exit... Hit CTRL-C to exit...
=====
...
[INFO ] launch-0!0:submitting job 1f (for 1) --> Job # 1f-0-1 <-depends-on None
[INFO ] launch-0!0:submitting job chk_1f (for 1) --> Job # chk_1f-0-1 <-depends-on_
↪1f-0-1
[INFO ] 1f-1!0:User error detected!!! for step 1f attempt 0 : (1)
[INFO ] chk_1f-1!0:User error detected!!! for step 1f attempt 0 : (1)
[INFO ] chk_1f-1!0:!!!!!!!!!! oooops pb : job missing or uncomplete at last step 1f :_
↪(1)
[INFO ] chk_1f-1!0:RESTARTING THE WRONG PART PREVIOUS JOB : 1f (1). current_
↪attempt=0 initial_attempt=0 Extra attempt #1 ( 1 out of 3) @ (2018-02-21 12:39:04)
[INFO ] chk_1f-1!0:submitting job 1f (for 1) --> Job # 1f-1-1 <-depends-on None
[INFO ] chk_1f-1!0:submitting job chk_1f (for 1) --> Job # chk_1f-0-1 <-depends-on_
↪219553
[INFO ] chk_1f-1!0:Job has been fixed and is restarting @ (2018-02-21 12:39:05)
[INFO ] 1f-1!1:User error detected!!! for step 1f attempt 1 : (1)
[INFO ] chk_1f-1!0:User error detected!!! for step 1f attempt 1 : (1)
[INFO ] chk_1f-1!0:!!!!!!!!!! oooops pb : job missing or uncomplete at last step 1f :_
↪(1)
[INFO ] chk_1f-1!0:RESTARTING THE WRONG PART PREVIOUS JOB : 1f (1). current_
↪attempt=1 initial_attempt=0 Extra attempt #2 ( 2 out of 3) @ (2018-02-21 12:39:18)
[INFO ] chk_1f-1!0:submitting job 1f (for 1) --> Job # 1f-2-1 <-depends-on None
[INFO ] chk_1f-1!0:submitting job chk_1f (for 1) --> Job # chk_1f-0-1 <-depends-on_
↪219555
[INFO ] chk_1f-1!0:Job has been fixed and is restarting @ (2018-02-21 12:39:19)
[INFO ] 1f-1!2:User error detected!!! for step 1f attempt 2 : (1)
[INFO ] chk_1f-1!0:User error detected!!! for step 1f attempt 2 : (1)
[INFO ] chk_1f-1!0:!!!!!!!!!! oooops pb : job missing or uncomplete at last step 1f :_
↪(1)
[INFO ] chk_1f-1!0:RESTARTING THE WRONG PART PREVIOUS JOB : 1f (1). current_
↪attempt=2 initial_attempt=0 Extra attempt #3 ( 3 out of 3) @ (2018-02-21 12:39:33)
[INFO ] chk_1f-1!0:submitting job 1f (for 1) --> Job # 1f-3-1 <-depends-on None
[INFO ] chk_1f-1!0:submitting job chk_1f (for 1) --> Job # chk_1f-0-1 <-depends-on_
↪219557
[INFO ] chk_1f-1!0:Job has been fixed and is restarting @ (2018-02-21 12:39:34)
[INFO ] 1f-1!3:User error detected!!! for step 1f attempt 3 : (1)
[INFO ] chk_1f-1!0:User error detected!!! for step 1f attempt 3 : (1)
[INFO ] chk_1f-1!0:!!!!!!!!!! oooops pb : job missing or uncomplete at last step 1f :_
↪(1)
[INFO ] chk_1f-1!0:RESTARTING THE WRONG PART PREVIOUS JOB : 1f (1). current_
↪attempt=3 initial_attempt=0 Extra attempt #4 ( 4 out of 3) @ (2018-02-21 12:39:46)
[INFO ] chk_1f-1!0:Too much failed attempt for step 1f my_joid is 219558 @ (2018-02-
↪21 12:39:46)
[INFO ] chk_1f-1!0:killing all the dependent jobs...
[INFO ] chk_1f-1!0:killing all the dependent jobs...
[INFO ] chk_1f-1!0:3 jobs to kill...
```

```
[INFO ] chk_1f-1!0:killing the job 219552 (step chk_1f-0)...
[INFO ] chk_1f-1!0:killing the job 219554 (step chk_1f-0)...
[INFO ] chk_1f-1!0:killing the job 219556 (step chk_1f-0)...
[INFO ] chk_1f-1!0:===== workflow is aborting =====
[INFO ] launch-0!0:===== workflow is finishing =====
```

## 5.4 Setting the number of restart

The faulty job is restarted automatically three times before *Decimate* declares the workflow as aborted. Restarting faulty job three times before aborting is the value set per default. It can be changed by adding `-max-retry=<your_value>` when submitting the job:

```
dbatch --max-retry=1 --check=check_job.sh --job-name=job_1 my_job_failed.sh
```

In this case *Decimate* only restarts the faulty job once, after two successive failed attempts:

```
dlog
()
=====
Currently Tailing ...
/home/kortass/DECIMATE-GITHUB/.decimate/LOGS/decimate.log
Hit CTRL-C to exit... Hit CTRL-C to exit...
=====
...
[INFO ] launch-0!0:submitting job 1f (for 1) --> Job # 1f-0-1 <-depends-on None
[INFO ] launch-0!0:submitting job chk_1f (for 1) --> Job # chk_1f-0-1 <-depends-on_
↪1f-0-1
[INFO ] 1f-1!0:User error detected!!! for step 1f attempt 0 : (1)
[INFO ] chk_1f-1!0:User error detected!!! for step 1f attempt 0 : (1)
[INFO ] chk_1f-1!0:!!!!!!!!!! oooops pb : job missing or uncomplete at last step 1f :_
↪(1)
[INFO ] chk_1f-1!0:RESTARTING THE WRONG PART PREVIOUS JOB : 1f (1). current_
↪attempt=0 initial_attempt=0 Extra attempt #1 ( 1 out of 1) @ (2018-02-21 12:44:53)
[INFO ] chk_1f-1!0:submitting job 1f (for 1) --> Job # 1f-1-1 <-depends-on None
[INFO ] chk_1f-1!0:submitting job chk_1f (for 1) --> Job # chk_1f-0-1 <-depends-on_
↪219561
[INFO ] chk_1f-1!0:Job has been fixed and is restarting @ (2018-02-21 12:44:54)
[INFO ] 1f-1!1:User error detected!!! for step 1f attempt 1 : (1)
[INFO ] chk_1f-1!0:User error detected!!! for step 1f attempt 1 : (1)
[INFO ] chk_1f-1!0:!!!!!!!!!! oooops pb : job missing or uncomplete at last step 1f :_
↪(1)
[INFO ] chk_1f-1!0:RESTARTING THE WRONG PART PREVIOUS JOB : 1f (1). current_
↪attempt=1 initial_attempt=0 Extra attempt #2 ( 2 out of 1) @ (2018-02-21 12:45:09)
[INFO ] chk_1f-1!0:Too much failed attempt for step 1f my_joid is 219562 @ (2018-02-
↪21 12:45:09)
[INFO ] chk_1f-1!0:killing all the dependent jobs...
[INFO ] chk_1f-1!0:killing all the dependent jobs...
[INFO ] chk_1f-1!0:1 jobs to kill...
[INFO ] chk_1f-1!0:killing the job 219560 (step chk_1f-0)...
[INFO ] chk_1f-1!0:===== workflow is aborting =====
[INFO ] launch-0!0:===== workflow is finishing =====
```

---

## Parameters combination

---

Then submission of parametric jobs requires to gather in a *parameter* file all the combinations of parameters that one wants to run a job against. This list of combination can be described as an explicit array of values of programatically via a Python or shell script or using simple directives.

While the execution of parametric workflows is described [here](#), here are detailed four ways of defining parameters. .

### 6.1 array of values

The simplest way to describe the set of parameter combinations that needs to be tested consists in listing them extensively as an array of values. The first row of this array is the name of each parameters and each row is one possible combination.

Here is a parameters file listing all possible combinations for 3 parameters (i,j,k), each of them taking the value 1 or 2.

```
# array-like description of parameter combinations
i j k
1 1 1
1 1 2
1 2 1
1 2 2
2 1 1
2 1 2
2 2 1
2 2 2
```

Notice that:

- spaces, void lines are ignored.
- every thing following a # is considered as a comment and ignored

## 6.2 Combined parameter sweep

In case of combinations that sweeps all possible set of values based on the domain definition of each variable, a more compact declarative syntax is also available. The same set of parameters can be generated with the following file:

```
# combine-like description of parameter combinations

#DECIM COMBINE i = [1,2]
#DECIM COMBINE j = [1,2]
#DECIM COMBINE k = [1,2]
```

Every line starting with `#DECIM` is parsed as a special command.

## 6.3 Parameters depending on simple formulas

Some parameters can also be computed from others using simple arithmetic formulas. Here is a way to declare them:

```
# combine-like description of parameter combinations

#DECIM COMBINE i = [1,2]
#DECIM COMBINE j = [1,2]
#DECIM COMBINE k = [1,2]

#DECIM p = i*j*k
```

which is a short way to describe the same 8 combinations as expressed in the following array-like parameter file:

```
# array-like description of parameter combinations

i j k p
1 1 1 1
1 1 2 2
1 2 1 2
1 2 2 4
2 1 1 2
2 1 2 4
2 2 1 4
2 2 2 8
```

an additional parameter can also be described by a list of values:

```
# combine-like description of parameter combinations

#DECIM COMBINE i = [1,2]
#DECIM COMBINE j = [1,2]
#DECIM COMBINE k = [1,2]

#DECIM p = i*j*k

#DECIM t = [1,2,4,8,16,32,64,128,256]
```

which is a short way to describe the same 8 combinations as expressed in the following array-like parameter file:

```
# array-like description of parameter combinations

i j k p t
1 1 1 1 1
1 1 2 2 2
1 2 1 2 4
1 2 2 4 8
2 1 1 2 16
2 1 2 4 32
2 2 1 4 64
2 2 2 8 128
```

For each parameter added via a list of values, the conformance with the existing number of already possible combinations is checked. For example, the following parameter file...

```
# combine-like description of parameter combinations

#DECIM COMBINE i = [1,2]
#DECIM COMBINE j = [1,2]
#DECIM COMBINE k = [1,2]

#DECIM p = i*j*k

#DECIM t = [1,2,4,8,16,32,64,128,256]
```

... produces the error:

```
[ERROR] parameters number mismatch for expression
[ERROR]      t = [1,2,4,8,16,32,64,128,256]
[ERROR]      --> expected 8 and got 9 parameters...
```

## 6.4 More complex Python expressions

For a high number of parameters, a portion of code written in Python can also be embedded after a `#DECIM PYTHON` directive till the end of the file.

```
# pythonic parameter example file

#DECIM COMBINE nodes = [2,4,8]
#DECIM COMBINE ntasks_per_node = [16,32]

#DECIM k = range(1,7)

#DECIM PYTHON

import math

ntasks = nodes*ntasks_per_node
nthreads = ntasks * 2

NPROC = 2; #Number of processors

t = int(2**(k))
T = 15
```

which is a short way to describe the same 8 combinations as expressed in the following array-like parameter file:

```
# array-like description of parameter combinations
nodes  ntasks_per_node  k  ntasks  nthreads  t  NPROC  T
  2           32  1     64     128    2     2    15
  2           64  2    128     256    4     2    15
  4           32  3    128     256    8     2    15
  4           64  4    256     512   16     2    15
  8           32  5    256     512   32     2    15
  8           64  6    512    1024   64     2    15
```

A python section is always evaluated at the end. Each new variables set at the end of the evaluation is added as a new parameter computed against each of the already built combinations. The conformance to the number of combinations already set is also checked if the variable is a set of values.

## 7.1 dbatch

Usage: dbatch [OPTIONS...] job\_script [args...]

Help:

- h, --help** show all possible options for **dbatch**
- H, --decimate-help** show hidden option to manage *Decimate* engine

Workflow management:

- kill** kill all jobs in the workflow either RUNNING, PENDING or WAITING
- resume** resume the already launched step and workflow in this directory
- restart** restart the already launched step or workflow in this directory
- ch, --check** check the step at its end (job DONE printed)

**-chf, --check-file=SCRIPT\_FILE** python or shell to check if results are ok

- xj, --max-jobs=MAX\_JOBS** maximum number of jobs to keep active in the queue (450 per default)
- xr, --max-retry=MAX\_RETRY** number of time a step can fail and be restarted automatically before failing the whole workflow (3 per default)

Execution in a pool:

- xy, --yalla** Execute simultaneous runs within a pool of nodes
- xyp, --yalla-parallel-runs=YALLA\_PARALLEL\_RUNS** number of parallel runs in a pool

Burst Buffer:

- bbz, --use-burst-buffer-size** use a non persistent burst buffer space
- xz, --burst-buffer-size=BURST\_BUFFER\_SIZE** set Burst Buffer space size
- bbs, --use-burst-buffer-space** use a persistent burst buffer space

`-xs, --burst-buffer-space=BURST_BUFFER_SPACE_name` sets Burst Buffer name

## 7.2 dstat

Usage: `dstat [OPTIONS...]`

Help:

`-h, --help` show all possible options for **dstat**

## 7.3 dlog

Usage: `dlog [OPTIONS...]`

Help:

`-h, --help` show all possible options for **dlog**

## 7.4 dkill

Usage: `dkill [OPTIONS...]`

Help:

`-h, --help` show all possible options for **dkill**

## 7.5 environment variables

environment variable forwarded to *Decimate* and setting option per default that will be added to any *Decimate* command initiated from the shell:

```
DPARAM
```

code to return when a job is considered as *Successful*:

```
0
```

code to return when a job is considered as *Failed*:

```
-1
```

code to return when a workflow has to be immediately stopped:

```
-9999
```

## 7.6 Job script directives

in script directives (to be added as-is anywhere in a SLURM job script).

To show the parameters set in the job environment from a parametic file processed via *Decimate*:

```
#DECIM SHOW_PARAMETERS
```

To process all the files ending by *.template* and replacing any parameter (typically *\_\_Name\_of\_parameter\_\_*) with a value coming from the parametric file processed by *Decimate*:

```
#DECIM PROCESS_TEMPLATE_FILES
```