
ReFrame Documentation

Release 2.21

CSCS

Dec 19, 2019

Table of Contents:

1	Use Cases	3
2	Latest Release	5
3	Publications	7
3.1	Getting Started	7
3.2	Configuring ReFrame for Your Site	9
3.3	The Regression Test Pipeline	17
3.4	ReFrame Tutorial	19
3.5	Customizing Further a Regression Test	39
3.6	How Test Dependencies Work In ReFrame	59
3.7	Understanding the Mechanism of Sanity Functions	61
3.8	Running ReFrame	69
3.9	Use Cases	92
3.10	About ReFrame	94
3.11	Reference Guide	95
3.12	Sanity Functions Reference	125
	Python Module Index	133
	Index	135

ReFrame is a new framework for writing regression tests for HPC systems. The goal of this framework is to abstract away the complexity of the interactions with the system, separating the logic of a regression test from the low-level details, which pertain to the system configuration and setup. This allows users to write easily portable regression tests, focusing only on the functionality.

Regression tests in ReFrame are simple Python classes that specify the basic parameters of the test. The framework will load the test and will send it down a well-defined pipeline that will take care of its execution. The stages of this pipeline take care of all the system interaction details, such as programming environment switching, compilation, job submission, job status query, sanity checking and performance assessment.

ReFrame also offers a high-level and flexible abstraction for writing sanity and performance checks for your regression tests, without having to care about the details of parsing output files, searching for patterns and testing against reference values for different systems.

Writing system regression tests in a high-level modern programming language, like Python, poses a great advantage in organizing and maintaining the tests. Users can create their own test hierarchies or test factories for generating multiple tests at the same time and they can also customize them in a simple and expressive way.

CHAPTER 1

Use Cases

The ReFrame framework has been in production at [CSCS](#) since the upgrade of the [Piz Daint](#) system in early December 2016.

[Read the full story...](#)

CHAPTER 2

Latest Release

ReFrame is being actively developed at [CSCS](#). You can always find the latest release [here](#).

- Slides [pdf] @ HPC System Testing BoF, SC'19.
- Slides [pdf] @ HUST 2019, SC'19.
- Slides [pdf] @ HPC Knowledge Meeting '19.
- Slides [pdf] & Talk @ FOSDEM'19.
- Slides [pdf] @ 4th EasyBuild User Meeting.
- Slides [pdf] @ HUST 2018, SC'18.
- Slides [pdf] [pptx] @ CSCS User Lab Day 2018.
- Slides [pdf] @ HPC Advisory Council 2018.
- Slides [pdf] @ SC17.
- Slides [pdf] @ CUG 2017.

3.1 Getting Started

3.1.1 Requirements

- Python 3.5 or higher. Python 2 is not supported.

Note: Changed in version 2.8: A functional TCL modules system is no more required. ReFrame can now operate without a modules system at all.

Optional

- For running the unit tests of the framework, the `pytest` unittesting framework is needed.

You are advised to run the *unit tests* of the framework after installing it on a new system to make sure that everything works fine.

3.1.2 Getting the Framework

To get the latest stable version of the framework, you can just clone it from the [github](#) project page:

```
git clone https://github.com/eth-cscs/reframe.git
```

Alternatively, you can pick a previous stable version by downloading it from the previous [releases](#) section.

3.1.3 Running the Unit Tests

After you have downloaded the framework, it is important to run the unit tests of to make sure that everything is set up correctly:

```
./test_reframe.py -v
```

The output should look like the following:

```
collected 442 items

unittests/test_argparser.py ..                               [  ]
↪0%]
unittests/test_cli.py .....s.....                          [  ]
↪4%]
unittests/test_config.py .....                             [  ]
↪7%]
unittests/test_deferrable.py .....                         [  ]
↪17%]
unittests/test_environments.py sss...s.....                 [  ]
↪20%]
unittests/test_exceptions.py .....                         [  ]
↪23%]
unittests/test_fields.py .....                             [  ]
↪28%]
unittests/test_launchers.py .....                          [  ]
↪31%]
unittests/test_loader.py .....                             [  ]
↪33%]
unittests/test_logging.py .....                            [  ]
↪38%]
unittests/test_modules.py .....ssssssssssssssss.....      [  ]
↪49%]
unittests/test_pipeline.py .....s..s.....                  [  ]
↪57%]
unittests/test_policies.py .....                           [  ]
↪64%]
unittests/test_runtime.py .                                 [  ]
↪64%]
unittests/test_sanity_functions.py .....                   [  ]
↪75%]
.....                                                       [  ]
↪78%]
unittests/test_schedulers.py .....s.s.....ss.....s.s.....ss. [  ]
↪90%]
```

(continues on next page)

(continued from previous page)

```

unittests/test_script_builders.py . [ ]
↪90%]
unittests/test_utility.py ..... [ ]
↪99%]
unittests/test_versioning.py .. [ ]
↪[100%]

===== 411 passed, 31 skipped in 28.10 seconds [ ]
↪=====

```

You will notice in the output that all the job submission related tests have been skipped. The test suite detects if the current system has a job submission system and is configured for ReFrame (see [Configuring ReFrame for your site](#)) and it will skip all the unsupported unit tests. As soon as you configure ReFrame for your system, you can rerun the test suite to check that job submission unit tests pass as well. Note here that some unit tests may still be skipped depending on the configured job submission system.

3.1.4 Where to Go from Here

The next step from here is to setup and configure ReFrame for your site, so that ReFrame can automatically recognize it and submit jobs. Please refer to the “[Configuring ReFrame For Your Site](#)” section on how to do that.

Before starting implementing a regression test, you should go through the “[The Regression Test Pipeline](#)” section, so as to understand the mechanism that ReFrame uses to run the regression tests. This section will let you follow easily the “[ReFrame Tutorial](#)” as well as understand the more advanced examples in the “[Customizing Further A Regression Test](#)” section.

To learn how to invoke the ReFrame command-line interface for running your tests, please refer to the “[Running ReFrame](#)” section.

3.2 Configuring ReFrame for Your Site

ReFrame provides an easy and flexible way to configure new systems and new programming environments. By default, it ships with a generic local system configured. This should be enough to let you run ReFrame on a local computer as soon as the basic [software requirements](#) are met.

As soon as a new system with its programming environments is configured, adapting an existing regression test could be as easy as just adding the system’s name in the `valid_systems` list and its associated programming environments in the `valid_prog_environs` list.

3.2.1 The Configuration File

The configuration of systems and programming environments is performed by a special Python dictionary called `site_configuration` defined inside the file `<install-dir>/reframe/settings.py`.

The `site_configuration` dictionary should define two entries, `systems` and `environments`. The former defines the systems that ReFrame may recognize, whereas the latter defines the available programming environments.

The following example shows a minimal configuration for the [Piz Daint](#) supercomputer at CSCS:

```

site_configuration = {
    'systems': {
        'daint': {

```

(continues on next page)

(continued from previous page)

```

'descr': 'Piz Daint',
'hostnames': ['daint'],
'modules_system': 'tmod',
'partitions': {
    'login': {
        'scheduler': 'local',
        'modules': [],
        'access': [],
        'environs': ['PrgEnv-cray', 'PrgEnv-gnu',
                    'PrgEnv-intel', 'PrgEnv-pgi'],
        'descr': 'Login nodes',
        'max_jobs': 4
    },
    'gpu': {
        'scheduler': 'nativeslurm',
        'modules': ['daint-gpu'],
        'access': ['--constraint=gpu'],
        'environs': ['PrgEnv-cray', 'PrgEnv-gnu',
                    'PrgEnv-intel', 'PrgEnv-pgi'],
        'container_platforms': {
            'Singularity': {
                'modules': ['Singularity']
            }
        },
        'descr': 'Hybrid nodes (Haswell/P100)',
        'max_jobs': 100
    },
    'mc': {
        'scheduler': 'nativeslurm',
        'modules': ['daint-mc'],
        'access': ['--constraint=mc'],
        'environs': ['PrgEnv-cray', 'PrgEnv-gnu',
                    'PrgEnv-intel', 'PrgEnv-pgi'],
        'container_platforms': {
            'Singularity': {
                'modules': ['Singularity']
            }
        },
        'descr': 'Multicore nodes (Broadwell)',
        'max_jobs': 100
    }
}
},
'environments': {
    '*': {
        'PrgEnv-cray': {
            'type': 'ProgEnvironment',
            'modules': ['PrgEnv-cray'],
        },
        'PrgEnv-gnu': {
            'type': 'ProgEnvironment',
            'modules': ['PrgEnv-gnu'],
        },
    }
}

```

(continues on next page)

(continued from previous page)

```

    },
    'PrgEnv-intel': {
        'type': 'ProgEnvironment',
        'modules': ['PrgEnv-intel'],
    },
    'PrgEnv-pgi': {
        'type': 'ProgEnvironment',
        'modules': ['PrgEnv-pgi'],
    }
}
}
}

```

3.2.2 System Configuration

The list of supported systems is defined as a set of key/value pairs under key `systems`. Each system is a key/value pair, with the key being the name of the system and the value being another set of key/value pairs defining its attributes. The valid attributes of a system are the following:

- `descr`: A detailed description of the system (default is the system name).
- `hostnames`: This is a list of hostname patterns according to the [Python Regular Expression Syntax](#), which will be used by ReFrame when it tries to *auto-detect* the current system (default []).
- `modules_system`: *[new in 2.8]* The modules system that should be used for loading environment modules on this system (default None). Three types of modules systems are currently supported:
 - `tmod` or `tmod32`: The classic Tcl implementation of the [environment modules](#) (version 3.2).
 - `tmod31`: *[new in 2.21]* The classic Tcl implementation of the [environment modules](#) (version 3.1).
 - `tmod4`: The version 4 of the Tcl implementation of the [environment modules](#) (versions older than 4.1 are not supported).
 - `lmod`: The Lua implementation of the [environment modules](#).
- `modules`: *[new in 2.19]* Modules to be loaded always when running on this system. These modules modify the ReFrame environment. This is useful when for example a particular module is needed to submit jobs on a specific system.
- `variables`: *[new in 2.19]* Environment variables to be set always when running on this system.
- `prefix`: Default regression prefix for this system (default `.`).
- `stagedir`: Default stage directory for this system (default None).
- `outputdir`: Default output directory for this system (default None).
- `perflogdir`: Default directory prefix for storing performance logs for this system (default None).
- `resourcesdir`: Default directory for storing large resources (e.g., input data files, etc.) needed by regression tests for this system (default `.`).
- `partitions`: A set of key/value pairs defining the partitions of this system and their properties (default {}). Partition configuration is discussed in the [next section](#).

For a more detailed description of the `prefix`, `stagedir`, `outputdir` and `perflogdir` directories, please refer to the “[Configuring ReFrame Directories](#)” and “[Performance Logging](#)” sections.

Note: A different backend is used for Tmod 3.1, due to its different Python bindings.

Warning: Changed in version 2.18: The `logdir` key is no more supported; please use `perflogdir` instead.

3.2.3 Partition Configuration

From the ReFrame's point of view, each system consists of a set of logical partitions. These partitions need not necessarily correspond to real scheduler partitions. For example, Piz Daint on the above example is split in *virtual partitions* using Slurm constraints. Other systems may be indeed split into real scheduler partitions.

The partitions of a system are defined similarly to systems as a set of key/value pairs with the key being the partition name and the value being another set of key/value pairs defining the partition's attributes. The available partition attributes are the following:

- `descr`: A detailed description of the partition (default is the partition name).
- `scheduler`: The job scheduler and parallel program launcher combination that is used on this partition to launch jobs. The syntax of this attribute is `<scheduler>+<launcher>`. A list of the supported *schedulers* and *parallel launchers* can be found at the end of this section.
- `access`: A list of scheduler options that will be passed to the generated job script for gaining access to that logical partition (default []).
- `environs`: A list of environments, with which ReFrame will try to run any regression tests written for this partition (default []). The environment names must be resolved inside the `environments` section of the `site_configuration` dictionary (see *Environments Configuration* for more information).
- `container_platforms`: *[new in 2.20]* A set of key/value pairs specifying the supported container platforms for this partition and how their environment is set up. Supported platform names are the following (names are case sensitive):
 - Docker: The *Docker* container runtime.
 - Singularity: The *Singularity* container runtime.
 - Sarus: The *Sarus* container runtime.

Each configured container runtime is associated optionally with an environment (modules and environment variables) that is providing it. This environment is specified as a dictionary in the following format:

```
{
  'modules': ['mod1', 'mod2', ...]
  'variables': {'ENV1': 'VAL1', 'ENV2': 'VAL2', ...}
}
```

If no special environment arrangement is needed for a configured container platform, you can simply specify an empty dictionary as an environment configuration, as it is shown in the following example:

```
'container_platforms': {
  'Docker': {}
}
```

- `modules`: A list of modules to be loaded before running a regression test on that partition (default []).
- `variables`: A set of environment variables to be set before running a regression test on that partition (default {}). Environment variables can be set as follows (notice that both the variable name and its value are strings):


```
'variables': {
  'MYVAR': '3',
  'OTHER': 'foo'
}
```

- `max_jobs`: The maximum number of concurrent regression tests that may be active (not completed) on this partition. This option is relevant only when ReFrame executes with the `asynchronous execution` policy.
- `resources`: A set of custom resource specifications and how these can be requested from the partition's scheduler (default `{}`).

This variable is a set of key/value pairs with the key being the resource name and the value being a list of options to be passed to the partition's job scheduler. The option strings can contain *placeholders* of the form `{placeholder_name}`. These placeholders may be replaced with concrete values by a regression tests through the `extra_resources` attribute.

For example, one could define a `gpu` resource for a multi-GPU system that uses Slurm as follows:

```
'resources': {
  'gpu': ['--gres=gpu:{num_gpus_per_node}']
}
```

A regression test then may request this resource as follows:

```
self.extra_resources = {'gpu': {'num_gpus_per_node': '8'}}
```

And the generated job script will have the following line in its preamble:

```
#SBATCH --gres=gpu:8
```

A resource specification may also start with `#PREFIX`, in which case `#PREFIX` will replace the standard job script prefix of the backend scheduler of this partition. This is useful in cases of job schedulers like Slurm, that allow alternative prefixes for certain features. An example is the `DataWarp` functionality of Slurm which is supported by the `#DW` prefix. One could then define `DataWarp` related resources as follows:

```
'resources': {
  'datawarp': [
    '#DW jobdw capacity={capacity} access_mode={mode} type=scratch',
    '#DW stage_out source={out_src} destination={out_dst} type={stage_
↵ filetype}'
  ]
}
```

A regression test that wants to make use of that resource, it can set its `extra_resources` as follows:

```
self.extra_resources = {
  'datawarp': {
    'capacity': '100GB',
    'mode': 'striped',
    'out_src': '$DW_JOB_STRIPED/name',
    'out_dst': '/my/file',
    'stage_filetype': 'file'
  }
}
```

Note: For the `PBS` backend, options accepted in the `access` and `resources` attributes may either refer to actual `qsub` options or be just resources specifications to be passed to the `-l select` option. The backend assumes a

qsub option, if the options passed in these attributes start with a `-`.

Note: Changed in version 2.8: A new syntax for the `scheduler` values was introduced as well as more parallel program launchers. The old values for the `scheduler` key will continue to be supported.

Note: Changed in version 2.9: Better support for custom job resources.

Note: Changed in version 2.14: The `modules` and `variables` partition configuration parameters do not affect the ReFrame environment anymore. They essentially define an environment to be always emitted when building and/or running the test on this partition. If you want to modify the environment ReFrame runs in for a particular system, define these parameters inside the *system configuration*.

Supported scheduler backends

ReFrame supports the following job schedulers:

- `slurm`: Jobs on the configured partition will be launched using `Slurm`. This scheduler relies on job accounting (`sacct` command) in order to reliably query the job status.
- `squeue`: *[new in 2.8.1]* Jobs on the configured partition will be launched using `Slurm`, but no job accounting is required. The job status is obtained using the `squeue` command. This scheduler is less reliable than the one based on the `sacct` command, but the framework does its best to query the job state as reliably as possible.
- `pbs`: *[new in 2.13]* Jobs on the configured partition will be launched using a `PBS-based` scheduler.
- `local`: Jobs on the configured partition will be launched locally as OS processes.

Supported parallel launchers

ReFrame supports the following parallel job launchers:

- `srun`: Programs on the configured partition will be launched using a bare `srun` command *without* any job allocation options passed to it. This launcher may only be used with the `slurm` scheduler.
- `srunalloc`: Programs on the configured partition will be launched using the `srun` command *with* job allocation options passed automatically to it. This launcher may also be used with the `local` scheduler.
- `alps`: Programs on the configured partition will be launched using the `aprun` command.
- `mpirun`: Programs on the configured partition will be launched using the `mpirun` command.
- `mpiexec`: Programs on the configured partition will be launched using the `mpiexec` command.
- `ibrun`: *[new in 2.21]* Programs on the configured partition will be launched using the `ibrun` command. This is a custom parallel job launcher used at `TACC`.
- `local`: Programs on the configured partition will be launched as-is without using any parallel program launcher.
- `ssh`: *[new in 2.20]* Programs on the configured partition will be launched using `SSH`. This option uses the partition's `access` parameter (see *above*) in order to determine the remote host and any additional options to be passed to the `SSH` client. The `ssh` command will be launched in “batch mode,” meaning that password-less access to the remote host must be configured. Here is an example configuration for the `ssh` launcher:

```
'partition_name': {
  'scheduler': 'local+ssh',
  'access': ['-l admin', 'remote.host'],
  'environs': ['builtin'],
}
```

Note that the environment is not propagated to the remote host, so the `environs` variable has no practical meaning except for enabling the testing of this partition.

There exist also the following aliases for specific combinations of job schedulers and parallel program launchers:

- `nativeslurm`: This is equivalent to `slurm+srund`.
- `local`: This is equivalent to `local+local`.

3.2.4 Environments Configuration

The environments available for testing in different systems are defined under the `environments` key of the top-level `site_configuration` dictionary. The `environments` key is associated to a special dictionary that defines scopes for looking up an environment. The `*` denotes the global scope and all environments defined there can be used by any system. Instead of `*`, you can define scopes for specific systems or specific partitions by using the name of the system or partition. For example, an entry `daint` will define a scope for a system called `daint`, whereas an entry `daint:gpu` will define a scope for a virtual partition named `gpu` on the system `daint`. When an environment name is used in the `environs` list of a system partition (see [Partition Configuration](#)), it is first looked up in the entry of that partition, e.g., `daint:gpu`. If no such entry exists, it is looked up in the entry of the system, e.g., `daint`. If not found there, it is looked up in the global scope denoted by the `*` key. If it cannot be found even there, an error will be issued. This look up mechanism allows you to redefine an environment for a specific system or partition. In the following example, we redefine `PrgEnv-gnu` for a system named `foo`, so that whenever `PrgEnv-gnu` is used on that system, the module `openmpi` will also be loaded and the compiler variables should point to the MPI wrappers.

```
'foo': {
  'PrgEnv-gnu': {
    'type': 'ProgEnvironment',
    'modules': ['PrgEnv-gnu', 'openmpi'],
    'cc': 'mpicc',
    'cxx': 'mpicxx',
    'ftn': 'mpif90',
  }
}
```

An environment is also defined as a set of key/value pairs with the key being its name and the value being a dictionary of its attributes. The possible attributes of an environment are the following:

- `type`: The type of the environment to create. There are two available environment types (note that names are case sensitive):
 - `'Environment'`: A simple environment.
 - `'ProgEnvironment'`: A programming environment.
- `modules`: A list of modules to be loaded when this environment is used (default `[]`, valid for all environment types)
- `variables`: A set of variables to be set when this environment is used (default `{}`, valid for all environment types)
- `cc`: The C compiler (default `'cc'`, valid for `'ProgEnvironment'` only).
- `cxx`: The C++ compiler (default `'CC'`, valid for `'ProgEnvironment'` only).

- `ftn`: The Fortran compiler (default `'ftn'`, valid for `'ProgEnvironment'` only).
- `cppflags`: The default preprocessor flags (default `None`, valid for `'ProgEnvironment'` only).
- `cflags`: The default C compiler flags (default `None`, valid for `'ProgEnvironment'` only).
- `cxxflags`: The default C++ compiler flags (default `None`, valid for `'ProgEnvironment'` only).
- `fflags`: The default Fortran compiler flags (default `None`, valid for `'ProgEnvironment'` only).
- `ldflags`: The default linker flags (default `None`, valid for `'ProgEnvironment'` only).

Note: All flags for programming environments are now defined as list of strings instead of simple strings.

Changed in version 2.17.

3.2.5 System Auto-Detection

When ReFrame is launched, it tries to detect the current system and select the correct site configuration entry. The auto-detection process is as follows:

ReFrame first tries to obtain the hostname from `/etc/xthostname`, which provides the unqualified *machine name* in Cray systems. If this cannot be found the hostname will be obtained from the standard `hostname` command. Having retrieved the hostname, ReFrame goes through all the systems in its configuration and tries to match the hostname against any of the patterns in the `hostnames` attribute of *system configuration*. The detection process stops at the first match found, and the system it belongs to is considered as the current system. If the system cannot be auto-detected, ReFrame will issue a warning and fall back to a generic system configuration, which is equivalent to the following:

```
site_configuration = {
  'systems': {
    'generic': {
      'descr': 'Generic fallback system configuration',
      'hostnames': ['localhost'],
      'partitions': {
        'login': {
          'scheduler': 'local',
          'environs': ['builtin-gcc'],
          'descr': 'Login nodes'
        }
      }
    }
  },
  'environments': {
    '*': {
      'builtin-gcc': {
        'type': 'ProgEnvironment',
        'cc': 'gcc',
        'cxx': 'g++',
        'ftn': 'gfortran',
      }
    }
  }
}
```

You can override completely the auto-detection process by specifying a system or a system partition with the `--system` option (e.g., `--system daint` or `--system daint:gpu`).

Note: Instead of issuing an error, ReFrame falls back to a generic system configuration in case system auto-detection fails.

Changed in version 2.19.

3.2.6 Viewing the current system configuration

New in version 2.16.

It is possible to ask ReFrame to print the configuration of the current system or the configuration of any programming environment defined for the current system. There are two command-line options for performing these operations:

- `--show-config`: This option shows the current system's configuration and exits. It can be combined with the `--system` option in order to show the configuration of another system.
- `--show-config-env ENV`: This option shows the configuration of the programming environment `ENV` and exits. The environment `ENV` must be defined for any of the partitions of the current system. This option can also be combined with `--system` in order to show the configuration of a programming environment defined for another system.

3.3 The Regression Test Pipeline

The backbone of the ReFrame regression framework is the regression test pipeline. This is a set of well defined phases that each regression test goes through during its lifetime. The figure below depicts this pipeline in detail.

Fig. 1: The regression test pipeline

A regression test starts its life after it has been instantiated by the framework. This is where all the basic information of the test is set. At this point, although it is initialized, the regression test is not yet *live*, meaning that it does not run yet. The framework will then go over all the loaded and initialized checks (we will talk about the loading and selection phases later), it will pick the next partition of the current system and the next programming environment for testing and will try to run the test. If the test supports the current system partition and the current programming environment, it will be run and it will go through all the following seven phases:

1. Setup
2. Compilation
3. Running
4. Sanity checking
5. Performance checking
6. Cleanup

A test may implement some of them as no-ops. As soon as the test is finished, its resources are cleaned up and the framework's environment is restored. ReFrame will try to repeat the same procedure on the same regression test using the next programming environment and the next system partition until no further environments and partitions are left to be tested. In the following we elaborate on each of the individual phases of the lifetime of a regression test.

3.3.1 0. The Initialization Phase

This phase is not part of the regression test pipeline as shown above, but it is quite important, since during this phase the test is loaded into memory and initialized. As we shall see in the “[Tutorial](#)” and in the “[Customizing Further A ReFrame Regression Test](#)” sections, this is the phase where the *specification* of a test is set. At this point the current system is already known and the test may be set up accordingly. If no further differentiation is needed depending on the system partition or the programming environment, the test could go through the whole pipeline performing all of its work without the need to override any of the other pipeline stages. In fact, this is perhaps the most common case for most of the regression tests.

3.3.2 1. The Setup Phase

A regression test is instantiated once by the framework and it is then copied each time a new system partition or programming environment is tried. This first phase of the regression pipeline serves the purpose of preparing the test to run on the specified partition and programming environment by performing a number of operations described below:

Set up and load the test’s environment

At this point the environment of the current partition, the current programming environment and any test’s specific environment will be loaded. For example, if the current partition requires `slurm`, the current programming environment is `PrgEnv-gnu` and the test requires also `cuda-toolkit`, this phase will be equivalent to the following:

```
module load slurm
module unload PrgEnv-cray
module load PrgEnv-gnu
module load cuda-toolkit
```

Note that the framework automatically detects conflicting modules and unloads them first. So the user need not to care about the existing environment at all. She only needs to specify what is needed by her test.

Setup the test’s paths

Each regression test is associated with a stage directory and an output directory. The stage directory will be the working directory of the test and all of its resources will be copied there before running. The output directory is the directory where some important output files of the test will be kept. By default these are the generated job script file, the standard output and standard error. The user can also specify additional files to be kept in the test’s specification. At this phase, all these directories are created.

Prepare a job for the test

At this point a *job descriptor* will be created for the test. A job descriptor in ReFrame is an abstraction of the job scheduler’s functionality relevant to the regression framework. It is responsible for submitting a job in a job queue and waiting for its completion. ReFrame supports two job scheduler backends that can be combined with several different parallel program launchers. For a complete list of the job scheduler/parallel launchers combinations, please refer to “[Partition Configuration](#)”.

3.3.3 2. The Compilation Phase

At this phase the source code associated with test is compiled with the current programming environment. Before compiling, all the resources of the test are copied to its stage directory and the compilation is performed from that

directory.

3.3.4 3. The Run Phase

This phase comprises two subphases:

- **Job launch:** At this subphase a job script file for the regression test is generated and submitted to the job scheduler queue. If the job scheduler for the current partition is the **local** one, a simple wrapper shell script will be generated and will be launched as a local OS process.
- **Job wait:** At this subphase the job (or local process) launched in the previous subphase is waited for. This phase is pretty basic: it just checks that the launched job (or local process) has finished. No check is made of whether the job or process has finished successfully or not. This is the responsibility of the next pipeline stage.

ReFrame currently supports two execution policies:

- **serial:** In the serial execution policy, these two subphases are performed back-to-back and the framework blocks until the current regression test finishes.
- **asynchronous:** In the asynchronous execution policy, as soon as the job associated to the current test is launched, ReFrame continues its execution by executing and launching the subsequent test cases.

3.3.5 4. The Sanity Checking Phase

At this phase it is determined whether the check has finished successfully or not. Although this decision is test-specific, ReFrame provides a very flexible and expressive way for specifying complex patterns and operations to be performed on the test's output in order to determine the outcome of the test.

3.3.6 5. The Performance Checking Phase

At this phase the performance of the regression test is checked. ReFrame uses the same mechanism for analyzing the output of the test as with sanity checking. The only difference is that the user can now specify reference values per system or system partition, as well as acceptable performance thresholds

3.3.7 6. The Cleanup Phase

This is the final stage of the regression test pipeline and it is responsible for cleaning up the resources of the test. Three steps are performed in this phase:

1. The interesting files of the test (job script, standard output and standard error and any additional files specified by the user) are copied to its output directory for later inspection and bookkeeping,
2. the stage directory is removed and
3. the test's environment is revoked.

At this point the ReFrame's environment is clean and in its original state and the framework may continue by running more test cases.

3.4 ReFrame Tutorial

This tutorial will guide you through writing your first regression tests with ReFrame. We will start with the most common and simple case of a regression test that compiles a code, runs it and checks its output. We will then expand

this example gradually by adding functionality and more advanced sanity and performance checks. By the end of the tutorial, you should be able to start writing your first regression tests with ReFrame.

If you just want to get a quick feeling of how it is like writing a regression test in ReFrame, you can start directly from here. However, if you want to get a better understanding of what is happening behind the scenes, we recommend to have a look also in “The Regression Test Pipeline” section.

All the tutorial examples can be found in `<reframe-install-prefix>/tutorial/`.

For the configuration of the system, we provide a minimal configuration file for Piz Daint, where we have tested all the tutorial examples. The site configuration that we used for this tutorial is the following:

```
'systems': {
  'daint': {
    'descr': 'Piz Daint',
    'hostnames': ['daint'],
    'modules_system': 'tmod',
    'partitions': {
      'login': {
        'scheduler': 'local',
        'modules': [],
        'access': [],
        'environs': ['PrgEnv-cray', 'PrgEnv-gnu',
                    'PrgEnv-intel', 'PrgEnv-pgi'],
        'descr': 'Login nodes',
        'max_jobs': 4
      },

      'gpu': {
        'scheduler': 'nativeslurm',
        'modules': ['daint-gpu'],
        'access': ['--constraint=gpu'],
        'environs': ['PrgEnv-cray', 'PrgEnv-gnu',
                    'PrgEnv-intel', 'PrgEnv-pgi'],
        'container_platforms': {
          'Singularity': {
            'modules': ['Singularity']
          }
        },
        'descr': 'Hybrid nodes (Haswell/P100)',
        'max_jobs': 100
      },

      'mc': {
        'scheduler': 'nativeslurm',
        'modules': ['daint-mc'],
        'access': ['--constraint=mc'],
        'environs': ['PrgEnv-cray', 'PrgEnv-gnu',
                    'PrgEnv-intel', 'PrgEnv-pgi'],
        'container_platforms': {
          'Singularity': {
            'modules': ['Singularity']
          }
        },
        'descr': 'Multicore nodes (Broadwell)',
        'max_jobs': 100
      }
    }
  }
}
```

(continues on next page)

(continued from previous page)

```

    },
    'environments': {
        '*': {
            'PrgEnv-cray': {
                'type': 'ProgEnvironment',
                'modules': ['PrgEnv-cray'],
            },
            'PrgEnv-gnu': {
                'type': 'ProgEnvironment',
                'modules': ['PrgEnv-gnu'],
            },
            'PrgEnv-intel': {
                'type': 'ProgEnvironment',
                'modules': ['PrgEnv-intel'],
            },
        },
    },
}

```

You can find the full `settings.py` file ready to be used by ReFrame in `<reframe-install-prefix>/tutorial/config/settings.py`. You may first need to go over the “Configuring ReFrame For Your Site” section, in order to prepare the framework for your systems.

3.4.1 The First Regression Test

The following is a simple regression test that compiles and runs a serial C program, which computes a matrix-vector product (`tutorial/src/example_matrix_multiplication.c`), and verifies its sane execution. As a sanity check, it simply looks for a specific output in the output of the program. Here is the full code for this test:

```

import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class Example1Test(rfm.RegressionTest):
    def __init__(self):
        self.descr = 'Simple matrix-vector multiplication example'
        self.valid_systems = ['*']
        self.valid_prog_environ = ['*']
        self.sourcepath = 'example_matrix_vector_multiplication.c'
        self.executable_opts = ['1024', '100']
        self.sanity_patterns = sn.assert_found(
            r'time for single matrix vector multiplication', self.stdout)
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}

```

A regression test written in ReFrame is essentially a Python class that must eventually derive from `RegressionTest`. To make a test visible to the framework, you must decorate your final test class with one of the following decorators:

- `@simple_test`: for registering a single parameterless instantiation of your test.
- `@parameterized_test`: for registering multiple instantiations of your test.

Let’s see in more detail how the `Example1Test` is defined:

```
@rfm.simple_test
class Example1Test(rfm.RegressionTest):
    def __init__(self):
```

The `__init__()` method is the constructor of your test. It is usually the only method you need to implement for your tests, especially if you don't want to customize any of the regression test pipeline stages. When your test is instantiated, the framework assigns a default name to it. This name is essentially a concatenation of the fully qualified name of the class and string representations of the constructor arguments, with any non-alphanumeric characters converted to underscores. In this example, the auto-generated test name is simply `Example1Test`. You may change the name of the test later in the constructor by setting the `name` attribute.

Note: Calling `super().__init__()` inside the constructor of a test is no more needed.

Changed in version 2.19.

Warning:

ReFrame requires that the names of all the tests it loads are unique. In case of name clashes, it will refuse to load the conflicting test.

New in version 2.12.

The next line sets a more detailed description of the test:

```
self.descr = 'Simple matrix-vector multiplication example'
```

This is optional and it defaults to the auto-generated test's name, if not specified.

Note: If you explicitly set only the name of the test, the description will not be automatically updated and will still keep its default value.

The next two lines specify the systems and the programming environments that this test is valid for:

```
self.valid_systems = ['*']
self.valid_prog_environs = ['*']
```

Both of these variables accept a list of system names or environment names, respectively. The `*` symbol is a wildcard meaning any system or any programming environment. The system and environment names listed in these variables must correspond to names of systems and environments defined in the ReFrame's [settings file](#).

When specifying system names you can always specify a partition name as well by appending `:<partname>` to the system's name. For example, given the configuration for our tutorial, `daint:gpu` would refer specifically to the `gpu` virtual partition of the system `daint`. If only a system name (without a partition) is specified in the `self.valid_systems` variable, e.g., `daint`, it means that this test is valid for any partition of this system.

The next line specifies the source file that needs to be compiled:

```
self.sourcepath = 'example_matrix_vector_multiplication.c'
```

ReFrame expects any source files, or generally resources, of the test to be inside an `src/` directory, which is at the same level as the regression test file. If you inspect the directory structure of the `tutorial/` folder, you will notice that:

```
tutorial/
  example1.py
  src/
    example_matrix_vector_multiplication.c
```

Notice also that you need not specify the programming language of the file you are asking ReFrame to compile or the compiler to use. ReFrame will automatically pick the correct compiler based on the extension of the source file. The exact compiler that is going to be used depends on the programming environment that the test is running with. For example, given our configuration, if it is run with `PrgEnv-cray`, the Cray C compiler will be used, if it is run with `PrgEnv-gnu`, the GCC compiler will be used etc. A user can associate compilers with programming environments in the ReFrame’s [settings file](#).

The next line in our first regression test specifies a list of options to be used for running the generated executable (the matrix dimension and the number of iterations in this particular example):

```
self.executable_opts = ['1024', '100']
```

Notice that you do not need to specify the executable name. Since ReFrame compiled it and generated it, it knows the name. We will see in the [“Customizing Further A ReFrame Regression Test”](#) section, how you can specify the name of the executable, in cases that ReFrame cannot guess its name.

The next lines specify what should be checked for assessing the sanity of the result of the test:

```
self.sanity_patterns = sn.assert_found(
    r'time for single matrix vector multiplication', self.stdout)
```

This expression simply asks ReFrame to look for `time for single matrix vector multiplication` in the standard output of the test. The `sanity_patterns` attribute can only be assigned the result of a special type of functions, called *sanity functions*. *Sanity functions* are special in the sense that they are evaluated lazily. You can generally treat them as normal Python functions inside a `sanity_patterns` expression. ReFrame provides already a wide range of useful sanity functions ranging from wrappers to the standard built-in functions of Python to functions related to parsing the output of a regression test. For a complete listing of the available functions, please have a look at the [“Sanity Functions Reference”](#).

In our example, the `assert_found` function accepts a regular expression pattern to be searched in a file and either returns `True` on success or raises a `SanityError` in case of failure with a descriptive message. This function uses internally the `re` module of the Python standard library, so it may accept the same [regular expression syntax](#). As a file argument, `assert_found` accepts any filename, which will be resolved against the stage directory of the test. You can also use the `stdout` and `stderr` attributes to reference the standard output and standard error, respectively.

Tip: You need not to care about handling exceptions, and error handling in general, inside your test. The framework will automatically abort the execution of the test, report the error and continue with the next test case.

The last two lines of the regression test are optional, but serve a good role in a production environment:

```
self.maintainers = ['you-can-type-your-email-here']
self.tags = {'tutorial'}
```

In the `maintainers` attribute you may store a list of people responsible for the maintenance of this test. In case of failure, this list will be printed in the failure summary.

The `tags` attribute is a set of tags that you can assign to this test. This is useful for categorizing the tests and helps in quickly selecting the tests of interest. More about test selection, you can find in the [“Running ReFrame”](#) section.

Note: The values assigned to the attributes of a `RegressionTest` are validated and if they don’t have the correct

type, an error will be issued by ReFrame. For a list of all the attributes and their types, please refer to the “Reference Guide”.

Running the Tutorial Examples

ReFrame offers a rich command-line interface that allows you to control several aspects of its executions. A more detailed description can be found in the “Running ReFrame” section. Here we will only show you how to run a specific tutorial test:

```
./bin/reframe -C tutorial/config/settings.py -c tutorial/example1.py -r
```

If everything is configured correctly for your system, you should get an output similar to the following:

```
Command line: ./bin/reframe -C tutorial/config/settings.py -c tutorial/example1.py -r
Reframe version: 2.13-dev0
Launched by user: XXX
Launched on host: daint104
Reframe paths
=====
    Check prefix      :
    Check search path : 'tutorial/example1.py'
    Stage dir prefix  : /current/working/dir/stage/
    Output dir prefix : /current/working/dir/output/
    Logging dir       : /current/working/dir/logs
[=====] Running 1 check(s)
[=====] Started on Fri May 18 13:19:12 2018

[-----] started processing Example1Test (Simple matrix-vector multiplication_
↳example)
[ RUN      ] Example1Test on daint:login using PrgEnv-cray
[      OK  ] Example1Test on daint:login using PrgEnv-cray
[ RUN      ] Example1Test on daint:login using PrgEnv-gnu
[      OK  ] Example1Test on daint:login using PrgEnv-gnu
[ RUN      ] Example1Test on daint:login using PrgEnv-intel
[      OK  ] Example1Test on daint:login using PrgEnv-intel
[ RUN      ] Example1Test on daint:login using PrgEnv-pgi
[      OK  ] Example1Test on daint:login using PrgEnv-pgi
[ RUN      ] Example1Test on daint:gpu using PrgEnv-cray
[      OK  ] Example1Test on daint:gpu using PrgEnv-cray
[ RUN      ] Example1Test on daint:gpu using PrgEnv-gnu
[      OK  ] Example1Test on daint:gpu using PrgEnv-gnu
[ RUN      ] Example1Test on daint:gpu using PrgEnv-intel
[      OK  ] Example1Test on daint:gpu using PrgEnv-intel
[ RUN      ] Example1Test on daint:gpu using PrgEnv-pgi
[      OK  ] Example1Test on daint:gpu using PrgEnv-pgi
[ RUN      ] Example1Test on daint:mc using PrgEnv-cray
[      OK  ] Example1Test on daint:mc using PrgEnv-cray
[ RUN      ] Example1Test on daint:mc using PrgEnv-gnu
[      OK  ] Example1Test on daint:mc using PrgEnv-gnu
[ RUN      ] Example1Test on daint:mc using PrgEnv-intel
[      OK  ] Example1Test on daint:mc using PrgEnv-intel
[ RUN      ] Example1Test on daint:mc using PrgEnv-pgi
[      OK  ] Example1Test on daint:mc using PrgEnv-pgi
[-----] finished processing Example1Test (Simple matrix-vector multiplication_
↳example)
```

(continues on next page)

(continued from previous page)

```
[ PASSED ] Ran 12 test case(s) from 1 check(s) (0 failure(s))
[=====] Finished on Fri May 18 13:20:17 2018
```

Notice how our regression test is run on every partition of the configured system and for every programming environment.

Now that you have got a first understanding of how a regression test is written in ReFrame, let's try to expand our example.

Inspecting the ReFrame Generated Files

As described in the [regression test pipeline](#) section, ReFrame generates several files during the execution of a test. When developing or debugging a regression test it is important to be able to locate them and inspect them.

As soon as the *setup* stage of the test is executed, a stage directory specific to this test is generated. All the required resources for the test are copied to this directory, and this will be the working directory for the compilation, running, sanity and performance checking phases. If the test is successful, this stage directory is removed, unless the `--keep-stage-files` option is passed in the command line. Before removing this directory, ReFrame copies the following files to a dedicated output directory for this test:

- The generated build script and its standard output and standard error. This allows you to inspect exactly how your test was compiled.
- The generated run script and its standard output and standard error. This allows you to inspect exactly how your test was run and verify that the sanity checking was correct.
- Any other user-specified files.

If a regression test fails, its stage directory will not be removed. This allows you to reproduce exactly what ReFrame was trying to perform and will help you debug the problem with your test.

Let's rerun our first example and instruct ReFrame to keep the stage directory of the test, so that we can inspect it.

```
./bin/reframe -C tutorial/config/settings.py -c tutorial/example1.py -r --keep-stage-
→files
```

ReFrame creates a stage directory for each test case using the following pattern:

```
$STAGEDIR_PREFIX/<system>/<partition>/<prog-environ>/<test-name>
```

Let's pick the test case for the `gpu` partition and the `PrgEnv-gnu` programming environment from our first test to inspect. The default `STAGEDIR_PREFIX` is `./stage`:

```
cd stage/daint/gpu/PrgEnv-gnu/Example1Test/
```

If you do a listing in this directory, you will see all the files contained in the `tutorial/src` directory, as well as the following files:

```
rfm_Example1Test_build.err  rfm_Example1Test_job.err
rfm_Example1Test_build.out  rfm_Example1Test_job.out
rfm_Example1Test_build.sh   rfm_Example1Test_job.sh
```

The `rfm_Example1Test_build.sh` is the generated build script and the `.out` and `.err` are the compilation's standard output and standard error. Here is the generated build script for our first test:

```
#!/bin/bash

_onerror()
{
    exitcode=$?
    echo "-reframe: command \"\$BASH_COMMAND\" failed (exit code: $exitcode)"
    exit $exitcode
}

trap _onerror ERR

module load daint-gpu
module unload PrgEnv-cray
module load PrgEnv-gnu
cc example_matrix_vector_multiplication.c -o ./Example1Test
```

Similarly, the `rfm_Example1Test_job.sh` is the generated job script and the `.out` and `.err` files are the corresponding standard output and standard error. The generated job script for the test case we are currently inspecting is the following:

```
#!/bin/bash -l
#SBATCH --job-name="rfm_Example1Test_job"
#SBATCH --time=0:10:0
#SBATCH --ntasks=1
#SBATCH --output=rfm_Example1Test_job.out
#SBATCH --error=rfm_Example1Test_job.err
#SBATCH --constraint=gpu
module load daint-gpu
module unload PrgEnv-cray
module load PrgEnv-gnu
srun ./Example1Test 1024 100
```

It is interesting to check here the generated job script for the `login` partition of the example system, which does not use a workload manager:

```
cat stage/daint/login/PrgEnv-gnu/Example1Test/rfm_Example1Test_job.sh
```

```
#!/bin/bash -l
module unload PrgEnv-cray
module load PrgEnv-gnu
./Example1Test 1024 100
```

This is one of the advantages in using ReFrame: You do not have to care about the system-level details of the target system that your test is running. Based on its configuration, ReFrame will generate the appropriate commands to run your test.

3.4.2 Customizing the Compilation Phase

In this example, we write a regression test to compile and run the OpenMP version of the matrix-vector product program, that we have shown before. The full code of this test follows:

```
import reframe as rfm
import reframe.utility.sanity as sn
```

(continues on next page)

(continued from previous page)

```

@rfm.simple_test
class Example2aTest(rfm.RegressionTest):
    def __init__(self):
        self.descr = 'Matrix-vector multiplication example with OpenMP'
        self.valid_systems = ['*']
        self.valid_prog_environs = ['PrgEnv-cray', 'PrgEnv-gnu',
                                     'PrgEnv-intel', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_openmp.c'
        self.build_system = 'SingleSource'
        self.executable_opts = ['1024', '100']
        self.variables = {
            'OMP_NUM_THREADS': '4'
        }
        self.sanity_patterns = sn.assert_found(
            r'time for single matrix vector multiplication', self.stdout)
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}

    @rfm.run_before('compile')
    def setflags(self):
        env = self.current_environ.name
        if env == 'PrgEnv-cray':
            self.build_system.cflags = ['-homp']
        elif env == 'PrgEnv-gnu':
            self.build_system.cflags = ['-fopenmp']
        elif env == 'PrgEnv-intel':
            self.build_system.cflags = ['-openmp']
        elif env == 'PrgEnv-pgi':
            self.build_system.cflags = ['-mp']

```

This example introduces two new concepts:

1. We need to set the `OMP_NUM_THREADS` environment variable, in order to specify the number of threads to use with our program.
2. We need to specify different flags for the different compilers provided by the programming environments we are testing. Notice also that we now restrict the validity of our test only to the programming environments that we know how to handle (see the `valid_prog_environs`).

To define environment variables to be set during the execution of a test, you should use the `variables` attribute of the `RegressionTest` class. This is a dictionary, whose keys are the names of the environment variables and whose values are the values of the environment variables. Notice that both the keys and the values must be strings.

From version 2.14, ReFrame manages compilation of tests through the concept of build systems. Any customization of the build process should go through a build system. For straightforward cases, as in our first example, where no customization is needed, ReFrame automatically picks the correct build system to build the code. In this example, however, we want to set the flags for compiling the OpenMP code. Assuming our test supported only GCC, we could simply add the following lines in the `__init__()` method of our test:

```

self.build_system = 'SingleSource'
self.build_system.cflags = ['-fopenmp']

```

The `SingleSource` build system that we use here supports the compilation of a single file only. Each build system type defines a set of variables that the user can set. Based on the selected build system, ReFrame will generate a build script that will be used for building the code. The generated build script can be found in the `stage` or the `output directory of the test`, along with the output of the compilation. This way, you may reproduce exactly what ReFrame does in case of any errors. More on the build systems feature can be found [here](#).

Getting back to our test, simply setting the `cflags` to `-fopenmp` globally in the test will make it fail for programming environments other than `PrgEnv-gnu`, since the OpenMP flags vary for the different compilers. Ideally, we need to set the `cflags` differently for each programming environment. To achieve this we need to define a method that will set the compilation flags based on the current programming environment (i.e., the environment that test currently runs with) and schedule it to run before the `compile` stage of the test pipeline as follows:

```
@rfm.run_before('compile')
def setflags(self):
    env = self.current_environ.name
    if env == 'PrgEnv-cray':
        self.build_system.cflags = ['-homp']
    elif env == 'PrgEnv-gnu':
        self.build_system.cflags = ['-fopenmp']
    elif env == 'PrgEnv-intel':
        self.build_system.cflags = ['-openmp']
    elif env == 'PrgEnv-pgi':
        self.build_system.cflags = ['-mp']
```

In this function we retrieve the current environment from the `current_environ` attribute, so we can then differentiate the build system's flags based on its name. Note that we cannot retrieve the current programming environment inside the test's constructor, since as described in in “[The Regression Test Pipeline](#)” section, it is during the setup phase that a regression test is prepared for a new system partition and a new programming environment. The second important thing in this function is the `@run_before('compile')` decorator. This decorator will attach this function to the `compile` stage of the pipeline and will execute it before entering this stage. There are six pipeline stages that are defined and can accept this type of hooks: `setup`, `compile`, `run`, `sanity`, `performance` and `cleanup`. Similarly, to the `@run_before` decorator, there is also the `@run_after`, which will run the decorated function after the specified pipeline stage. The decorated function may have any name, but it should be a method of the test taking no arguments (i.e., its sole argument should be `self`).

You may attach multiple functions to the same stage as in the following example:

```
@rfm.run_before('compile')
def setflags(self):
    env = self.current_environ.name
    if env == 'PrgEnv-cray':
        self.build_system.cflags = ['-homp']
    elif env == 'PrgEnv-gnu':
        self.build_system.cflags = ['-fopenmp']
    elif env == 'PrgEnv-intel':
        self.build_system.cflags = ['-openmp']
    elif env == 'PrgEnv-pgi':
        self.build_system.cflags = ['-mp']

@rfm.run_before('compile')
def set_more_flags(self):
    self.build_system.flags += ['-g']
```

In this case, the decorated functions will be executed before the compilation stage in the order that they are defined in the regression test.

There is also the possibility to attach a single function to multiple stages by stacking the `@run_before` or `@run_after` decorators. In the following example `var` will be set to 2 after the setup phase is executed:

```
def __init__(self):
    ...
    self.var = 0
```

(continues on next page)

(continued from previous page)

```

@rfm.run_before('setup')
@rfm.run_after('setup')
def inc(self):
    self.var += 1

```

Another important feature of the hooks syntax, is that hooks are inherited by derived tests, unless you override the function and re-hook it explicitly. In the following example, the `setflags()` will be executed before the compilation phase of the `DerivedTest`:

```

class BaseTest(rfm.RegressionTest):
    def __init__(self):
        ...
        self.build_system = 'Make'

    @rfm.run_before('compile')
    def setflags(self):
        if self.current_envIRON.name == 'X':
            self.build_system.cppflags = ['-Ifoo']

@rfm.simple_test
class DerivedTest(BaseTest):
    def __init__(self):
        super().__init__()
        ...

```

If you override a hooked function in a derived class, the base class' hook will not be executed, unless you explicitly call it with `super()`. In the following example, we completely disable the `setflags()` hook of the base class:

```

@rfm.simple_test
class DerivedTest(BaseTest):
    @rfm.run_before('compile')
    def setflags(self):
        pass

```

Notice that in order to redefine a hook, you need not only redefine the method in the derived class, but you should hook it at the same pipeline phase. Otherwise, the base class hook will be executed.

Note: You may still configure your test per programming environment and per system partition by overriding the `setup` method, as in ReFrame versions prior to 2.20, but this is now discouraged since it is more error prone, as you have to memorize the signature of the pipeline methods that you override and also remember to call `super()`.

Warning: Setting the compiler flags in the programming environment has been dropped completely in version 2.17.

An alternative implementation using dictionaries

Here we present an alternative implementation of the same test using a dictionary to hold the compilation flags for the different programming environments. The advantage of this implementation is that you move the different compilation flags in the initialization phase, where also the rest of the test's specification is, thus making it more concise.

The `setup()` method is now very simple: it gets the correct compilation flags from the `prgenv_flags` dictionary and applies them to the build system.

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class Example2bTest(rfm.RegressionTest):
    def __init__(self):
        self.descr = 'Matrix-vector multiplication example with OpenMP'
        self.valid_systems = ['*']
        self.valid_prog_environ = ['PrgEnv-cray', 'PrgEnv-gnu',
                                   'PrgEnv-intel', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_openmp.c'
        self.build_system = 'SingleSource'
        self.executable_opts = ['1024', '100']
        self.prgenv_flags = {
            'PrgEnv-cray': ['-homp'],
            'PrgEnv-gnu': ['-fopenmp'],
            'PrgEnv-intel': ['-openmp'],
            'PrgEnv-pgi': ['-mp']
        }
        self.variables = {
            'OMP_NUM_THREADS': '4'
        }
        self.sanity_patterns = sn.assert_found(
            r'time for single matrix vector multiplication', self.stdout)
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}

    @rfm.run_before('compile')
    def setflags(self):
        self.build_system.cflags = self.prgenv_flags[self.current_envIRON.name]
```

Tip: A regression test is like any other Python class, so you can freely define your own attributes. If you accidentally try to write on a reserved `RegressionTest` attribute that is not writeable, ReFrame will prevent this and it will throw an error.

3.4.3 Running on Multiple Nodes

So far, all our tests run on a single node. Depending on the actual system that ReFrame is running, the test may run locally or be submitted to the system's job scheduler. In this example, we write a regression test for the MPI+OpenMP version of the matrix-vector product. The source code of this program is in `tutorial/src/example_matrix_vector_multiplication_mpi_openmp.c`. The regression test file follows:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class Example3Test(rfm.RegressionTest):
    def __init__(self):
        self.descr = 'Matrix-vector multiplication example with MPI'
```

(continues on next page)

(continued from previous page)

```

self.valid_systems = ['daint:gpu', 'daint:mc']
self.valid_prog_environ = ['PrgEnv-cray', 'PrgEnv-gnu',
                           'PrgEnv-intel', 'PrgEnv-pgi']
self.sourcepath = 'example_matrix_vector_multiplication_mpi_openmp.c'
self.executable_opts = ['1024', '10']
self.build_system = 'SingleSource'
self.prgenv_flags = {
    'PrgEnv-cray': ['-homp'],
    'PrgEnv-gnu': ['-fopenmp'],
    'PrgEnv-intel': ['-openmp'],
    'PrgEnv-pgi': ['-mp']
}
self.sanity_patterns = sn.assert_found(
    r'time for single matrix vector multiplication', self.stdout)
self.num_tasks = 8
self.num_tasks_per_node = 2
self.num_cpus_per_task = 4
self.variables = {
    'OMP_NUM_THREADS': str(self.num_cpus_per_task)
}
self.maintainers = ['you-can-type-your-email-here']
self.tags = {'tutorial'}

@rfm.run_before('compile')
def setflags(self):
    self.build_system.cflags = self.prgenv_flags[self.current_envIRON.name]

```

This test is pretty much similar to the *test example* for the OpenMP code we have shown before, except that it adds some information about the configuration of the distributed tasks. It also restricts the valid systems only to those that support distributed execution. Let's take the changes step-by-step:

First we need to specify for which partitions this test is meaningful by setting the *valid_systems* attribute:

```
self.valid_systems = ['daint:gpu', 'daint:mc']
```

We only specify the partitions that are configured with a job scheduler. If we try to run the generated executable on the login nodes, it will fail. So we remove this partition from the list of the supported systems.

The most important addition to this check are the variables controlling the distributed execution:

```
self.num_tasks = 8
self.num_tasks_per_node = 2
self.num_cpus_per_task = 4
```

By setting these variables, we specify that this test should run with 8 MPI tasks in total, using two tasks per node. Each task may use four logical CPUs. Based on these variables ReFrame will generate the appropriate scheduler flags to meet that requirement. For example, for Slurm these variables will result in the following flags: `--ntasks=8`, `--ntasks-per-node=2` and `--cpus-per-task=4`. ReFrame provides several more variables for configuring the job submission. As shown in the following Table, they follow closely the corresponding Slurm options. For schedulers that do not provide the same functionality, some of the variables may be ignored.

RegressionTest attribute	Corresponding SLURM option
<code>time_limit = (0, 10, 30)</code>	<code>--time=00:10:30</code>
<code>use_multithreading = True</code>	<code>--hint=multithread</code>
<code>use_multithreading = False</code>	<code>--hint=nomultithread</code>
<code>exclusive_access = True</code>	<code>--exclusive</code>
<code>num_tasks=72</code>	<code>--ntasks=72</code>
<code>num_tasks_per_node=36</code>	<code>--ntasks-per-node=36</code>
<code>num_cpus_per_task=4</code>	<code>--cpus-per-task=4</code>
<code>num_tasks_per_core=2</code>	<code>--ntasks-per-core=2</code>
<code>num_tasks_per_socket=36</code>	<code>--ntasks-per-socket=36</code>

3.4.4 Testing a GPU Code

In this example, we will create two regression tests for two different GPU versions of our matrix-vector code: OpenACC and CUDA. Let's start with the OpenACC regression test:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class Example4Test(rfm.RegressionTest):
    def __init__(self):
        self.descr = 'Matrix-vector multiplication example with OpenACC'
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environ = ['PrgEnv-cray', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_openacc.c'
        self.build_system = 'SingleSource'
        self.executable_opts = ['1024', '100']
        self.modules = ['craype-accel-nvidia60']
        self.num_gpus_per_node = 1
        self.prgenv_flags = {
            'PrgEnv-cray': ['-hacc', '-hnoomp'],
            'PrgEnv-pgi': ['-acc', '-ta=tesla:cc60']
        }
        self.sanity_patterns = sn.assert_found(
            r'time for single matrix vector multiplication', self.stdout)
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}

    @rfm.run_before('compile')
    def setflags(self):
        self.build_system.cflags = self.prgenv_flags[self.current_environs.name]
```

The things to notice in this test are the restricted list of system partitions and programming environments that this test supports and the use of the `modules` variable:

```
self.modules = ['craype-accel-nvidia60']
```

The `modules` variable takes a list of modules that should be loaded during the setup phase of the test. In this particular test, we need to load the `craype-accel-nvidia60` module, which enables the generation of a GPU binary from an OpenACC code.

It is also important to note that in GPU-enabled tests the number of GPUs for each node have to be specified by setting the corresponding variable `num_gpus_per_node`, as follows:

```
self.num_gpus_per_node = 1
```

The regression test for the CUDA code is slightly simpler:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class Example5Test(rfm.RegressionTest):
    def __init__(self):
        self.descr = 'Matrix-vector multiplication example with CUDA'
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environs = ['PrgEnv-cray', 'PrgEnv-gnu', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_cuda.cu'
        self.executable_opts = ['1024', '100']
        self.modules = ['cudatoolkit']
        self.num_gpus_per_node = 1
        self.sanity_patterns = sn.assert_found(
            r'time for single matrix vector multiplication', self.stdout)
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}
```

ReFrame will recognize the `.cu` extension of the source file and it will try to invoke `nvcc` for compiling the code. In this case, there is no need to differentiate across the programming environments, since the compiler will be eventually the same. `nvcc` in our example is provided by the `cudatoolkit` module, which we list it in the `modules` variable.

3.4.5 More Advanced Sanity Checking

So far we have done a very simple sanity checking. We are only looking if a specific line is present in the output of the test program. In this example, we expand the regression test of the serial code, so as to check also if the printed norm of the result vector is correct.

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class Example6Test(rfm.RegressionTest):
    def __init__(self):
        self.descr = 'Matrix-vector multiplication with L2 norm check'
        self.valid_systems = ['*']
        self.valid_prog_environs = ['*']
        self.sourcepath = 'example_matrix_vector_multiplication.c'

        matrix_dim = 1024
        iterations = 100
        self.executable_opts = [str(matrix_dim), str(iterations)]

        expected_norm = matrix_dim
        found_norm = sn.extractsingle(
            r'The L2 norm of the resulting vector is:\s+(?P<norm>\S+)',
            self.stdout, 'norm', float)
        self.sanity_patterns = sn.all([
            sn.assert_found(
                r'time for single matrix vector multiplication', self.stdout),
```

(continues on next page)

(continued from previous page)

```

        sn.assert_lt(sn.abs(expected_norm - found_norm), 1.0e-6)
    ])
    self.maintainers = ['you-can-type-your-email-here']
    self.tags = {'tutorial'}

```

The only difference with our first example is actually the more complex expression to assess the sanity of the test. Let's go over it line-by-line. The first thing we do is to extract the norm printed in the standard output.

```

found_norm = sn.extractsingle(
    r'The L2 norm of the resulting vector is:\s+(?P<norm>\S+)',
    self.stdout, 'norm', float)

```

The *extractsingle* sanity function extracts some information from a single occurrence (by default the first) of a pattern in a filename. In our case, this function will extract the `norm` capturing group from the match of the regular expression `r'The L2 norm of the resulting vector is:\s+(?P<norm>\S+)'` in standard output, it will convert it to float and it will return it. Unnamed capturing groups in regular expressions are also supported, which you can reference by their group number. For example, we could have written the same statement as follows:

```

found_norm = sn.extractsingle(
    r'The L2 norm of the resulting vector is:\s+(\S+)',
    self.stdout, 1, float)

```

Notice that we replaced the `'norm'` argument with `1`, which is the capturing group number.

Note: In regular expressions, capturing group 0 corresponds always to the whole match. In sanity functions dealing with regular expressions, this will yield the whole line that matched.

A useful counterpart of *extractsingle* is the *extractall* function, which instead of a single occurrence, returns a list of all the occurrences found. For a more detailed description of this and other sanity functions, please refer to the [sanity function reference](#).

The next four lines is the actual sanity check:

```

self.sanity_patterns = sn.all([
    sn.assert_found(
        r'time for single matrix vector multiplication', self.stdout),
    sn.assert_lt(sn.abs(expected_norm - found_norm), 1.0e-6)
])

```

This expression combines two conditions that need to be true, in order for the sanity check to succeed:

1. Find in standard output the same line we were looking for already in the first example.
2. Verify that the printed norm does not deviate significantly from the expected value.

The *all* function is responsible for combining the results of the individual subexpressions. It is essentially the Python built-in `all()` function, exposed as a sanity function, and requires that all the elements of the iterable it takes as an argument evaluate to `True`. As mentioned before, all the `assert_*` functions either return `True` on success or raise `SanityError`. So, if everything goes smoothly, `sn.all()` will evaluate to `True` and sanity checking will succeed.

The expression for the second condition is more interesting. Here, we want to assert that the absolute value of the difference between the expected and the found norm are below a certain value. The important thing to mention here is that you can combine the results of sanity functions in arbitrary expressions, use them as arguments to other functions, return them from functions, assign them to variables etc. Remember that sanity functions are not evaluated at the time you call them. They will be evaluated later by the framework during the sanity checking phase. If you include the result

of a sanity function in an expression, the evaluation of the resulting expression will also be deferred. For a detailed description of the mechanism behind the sanity functions, please have a look at “[Understanding The Mechanism Of Sanity Functions](#)” section.

3.4.6 Writing a Performance Test

An important aspect of regression testing is checking for performance regressions. ReFrame offers a flexible way of extracting and manipulating performance data from the program output, as well as a comprehensive way of setting performance thresholds per system and system partitions.

In this example, we extend the CUDA test presented [previously](#), so as to check also the performance of the matrix-vector multiplication.

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class Example7Test(rfm.RegressionTest):
    def __init__(self):
        self.descr = 'Matrix-vector multiplication (CUDA performance test)'
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environ = ['PrgEnv-gnu', 'PrgEnv-cray', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_cuda.cu'
        self.build_system = 'SingleSource'
        self.build_system.cxxflags = ['-O3']
        self.executable_opts = ['4096', '1000']
        self.modules = ['cudatoolkit']
        self.num_gpus_per_node = 1
        self.sanity_patterns = sn.assert_found(
            r'time for single matrix vector multiplication', self.stdout)
        self.perf_patterns = {
            'perf': sn.extractsingle(r'Performance:\s+(?P<Gflops>\S+) Gflop/s',
                                   self.stdout, 'Gflops', float)
        }
        self.reference = {
            'daint:gpu': {
                'perf': (50.0, -0.1, 0.1, 'Gflop/s'),
            }
        }
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}
```

There are two new variables set in this test that basically enable the performance testing:

perf_patterns This variable defines which are the performance patterns we are looking for and how to extract the performance values.

reference This variable is a collection of reference values for different systems.

Let’s have a closer look at each of them:

```
self.perf_patterns = {
    'perf': sn.extractsingle(r'Performance:\s+(?P<Gflops>\S+) Gflop/s',
                           self.stdout, 'Gflops', float)
}
```

The `perf_patterns` attribute is a dictionary, whose keys are *performance variables* (i.e., arbitrary names assigned to the performance values we are looking for), and its values are *sanity expressions* that specify how to obtain these performance values from the output. A sanity expression is a Python expression that uses the result of one or more *sanity functions*. In our example, we name the performance value we are looking for simply as `perf` and we extract its value by converting to float the regex capturing group named `Gflops` from the line that was matched in the standard output.

Each of the performance variables defined in `perf_patterns` must be resolved in the `reference` dictionary of reference values. When the framework obtains a performance value from the output of the test it searches for a reference value in the `reference` dictionary, and then it checks whether the user supplied tolerance is respected. Let's go over the `reference` dictionary of our example and explain its syntax in more detail:

```
self.reference = {
    'daint:gpu': {
        'perf': (50.0, -0.1, 0.1, 'Gflop/s'),
    }
}
```

This is a special type of dictionary that we call *scoped dictionary*, because it defines scopes for its keys. We have already seen it being used in the `environments` section of the `configuration file` of ReFrame. In order to resolve a reference value for a performance variable, ReFrame creates the following key `<current_sys>:<current_part>:<perf_variable>` and looks it up inside the `reference` dictionary. In our example, since this test is only allowed to run on the `daint:gpu` partition of our system, ReFrame will look for the `daint:gpu:perf` reference key. The `perf` subkey will then be searched in the following scopes in this order: `daint:gpu`, `daint`, `*`. The first occurrence will be used as the reference value of the `perf` performance variable. In our example, the `perf` key will be resolved in the `daint:gpu` scope giving us the reference value.

Reference values in ReFrame are specified as a three-tuple or four-tuple comprising the reference value, the lower and upper thresholds and, optionally, the measurement unit. Thresholds are specified as decimal fractions of the reference value. For nonnegative reference values, the lower threshold must lie in the `[-1,0]`, whereas the upper threshold may be any positive real number or zero. In our example, the reference value for this test on `daint:gpu` is 50 Gflop/s $\pm 10\%$. Setting a threshold value to `None` disables the threshold. If you specify a measurement unit as well, you will be able to log it the performance logs of the test; this is handy when you are inspecting or plotting the performance values.

ReFrame will always add a default `*` entry in the `reference` dictionary, if it does not exist, with the reference value of `(0, None, None, <unit>)`, where `unit` is derived from the unit of each respective performance variable. This is useful when using ReFrame for benchmarking purposes and you would like to run a test on an unknown system.

Note: Reference tuples may now optionally contain units.

New in version 2.16.

Note: A default `*` entry is now always added to the reference dictionary.

New in version 2.19.

3.4.7 Combining It All Together

As we have mentioned before and as you have already experienced with the examples in this tutorial, regression tests in ReFrame are written in pure Python. As a result, you can leverage the language features and capabilities to organize better your tests and decrease the maintenance cost. In this example, we are going to reimplement all the tests of the

tutorial with much less code and in a single file. Here is the final example code that combines all the tests discussed before:

```
import reframe as rfm
import reframe.utility.sanity as sn

class BaseMatrixVectorTest (rfm.RegressionTest):
    def __init__(self, test_version):
        self.descr = '%s matrix-vector multiplication' % test_version
        self.valid_systems = ['*']
        self.valid_prog_environ = ['*']
        self.build_system = 'SingleSource'
        self.prgenv_flags = None

        matrix_dim = 1024
        iterations = 100
        self.executable_opts = [str(matrix_dim), str(iterations)]

        expected_norm = matrix_dim
        found_norm = sn.extractsingle(
            r'The L2 norm of the resulting vector is:\s+(?P<norm>\S+)',
            self.stdout, 'norm', float)
        self.sanity_patterns = sn.all([
            sn.assert_found(
                r'time for single matrix vector multiplication', self.stdout),
            sn.assert_lt(sn.abs(expected_norm - found_norm), 1.0e-6)
        ])
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}

    @rfm.run_before('compile')
    def setflags(self):
        if self.prgenv_flags is not None:
            env = self.current_envIRON.name
            self.build_system.cflags = self.prgenv_flags[env]

@rfm.simple_test
class SerialTest (BaseMatrixVectorTest):
    def __init__(self):
        super().__init__('Serial')
        self.sourcepath = 'example_matrix_vector_multiplication.c'

@rfm.simple_test
class OpenMPTest (BaseMatrixVectorTest):
    def __init__(self):
        super().__init__('OpenMP')
        self.sourcepath = 'example_matrix_vector_multiplication_openmp.c'
        self.valid_prog_environ = ['PrgEnv-cray', 'PrgEnv-gnu',
                                    'PrgEnv-intel', 'PrgEnv-pgi']

        self.prgenv_flags = {
            'PrgEnv-cray': ['-homp'],
            'PrgEnv-gnu': ['-fopenmp'],
            'PrgEnv-intel': ['-openmp'],
            'PrgEnv-pgi': ['-mp']
        }
}
```

(continues on next page)

```

        self.variables = {
            'OMP_NUM_THREADS': '4'
        }

@rfm.simple_test
class MPITest(BaseMatrixVectorTest):
    def __init__(self):
        super().__init__('MPI')
        self.valid_systems = ['daint:gpu', 'daint:mc']
        self.valid_prog_environ = ['PrgEnv-cray', 'PrgEnv-gnu',
                                   'PrgEnv-intel', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_mpi_openmp.c'
        self.prgenv_flags = {
            'PrgEnv-cray': ['-homp'],
            'PrgEnv-gnu': ['-fopenmp'],
            'PrgEnv-intel': ['-openmp'],
            'PrgEnv-pgi': ['-mp']
        }
        self.num_tasks = 8
        self.num_tasks_per_node = 2
        self.num_cpus_per_task = 4
        self.variables = {
            'OMP_NUM_THREADS': str(self.num_cpus_per_task)
        }

@rfm.simple_test
class OpenACCTest(BaseMatrixVectorTest):
    def __init__(self):
        super().__init__('OpenACC')
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environ = ['PrgEnv-cray', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_openacc.c'
        self.modules = ['craype-accel-nvidia60']
        self.num_gpus_per_node = 1
        self.prgenv_flags = {
            'PrgEnv-cray': ['-hacc', '-hnoomp'],
            'PrgEnv-pgi': ['-acc', '-ta=tesla:cc60']
        }

@rfm.simple_test
class CudaTest(BaseMatrixVectorTest):
    def __init__(self):
        super().__init__('CUDA')
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environ = ['PrgEnv-gnu', 'PrgEnv-cray', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_cuda.cu'
        self.modules = ['cudatoolkit']
        self.num_gpus_per_node = 1

```

This test abstracts away the common functionality found in almost all of our tutorial tests (executable options, sanity checking, etc.) to a base class, from which all the concrete regression tests derive. Each test then redefines only the parts that are specific to it. Notice also that only the actual tests, i.e., the derived classes, are made visible to the framework through the `@simple_test` decorator. Decorating the base class has no meaning, because it does not correspond to an actual test.

The total line count of this refactored example is less than half of that of the individual tutorial tests. Another interesting thing to note here is the base class accepting additional additional parameters to its constructor, so that the concrete subclasses can initialize it based on their needs.

3.4.8 Summary

This concludes our ReFrame tutorial. We have covered all basic aspects of writing regression tests in ReFrame and you should now be able to start experimenting by writing your first useful tests. The [next section](#) covers further topics in customizing a regression test to your needs.

3.5 Customizing Further a Regression Test

In this section, we are going to show some more elaborate use cases of ReFrame. Through the use of more advanced examples, we will demonstrate further customization options which modify the default options of the ReFrame pipeline. The corresponding scripts as well as the source code of the examples discussed here can be found in the directory `tutorial/advanced`.

3.5.1 Working with Makefiles

We have already shown how you can compile a single source file associated with your regression test. In this example, we show how ReFrame can leverage Makefiles to build executables.

Compiling a regression test through a Makefile is straightforward with ReFrame. If the `sourcepath` attribute refers to a directory, then ReFrame will automatically invoke `make` in that directory. More specifically, ReFrame first copies the `sourcesdir` to the stage directory at the beginning of the compilation phase and then constructs the path `os.path.join('{STAGEDIR}', self.sourcepath)` to determine the actual compilation path. If this is a directory, it will invoke `make` in it.

Note: The `sourcepath` attribute must be a relative path referring to a subdirectory of `sourcesdir`, i.e., relative paths starting with `..` will be rejected.

By default, `sourcepath` is the empty string and `sourcesdir` is set to `'src/'`. As a result, by not specifying a `sourcepath` at all, ReFrame will eventually compile the files found in the `src/` directory. This is exactly what our first example here does.

For completeness, here are the contents of Makefile provided:

```
EXECUTABLE := advanced_example1

.SUFFIXES: .o .c

OBJS := advanced_example1.o

$(EXECUTABLE) : $(OBJS)
    $(CC) $(CFLAGS) $(LDFLAGS) -o $@ $^

$(OBJS) : advanced_example1.c
    $(CC) $(CPPFLAGS) $(CFLAGS) -c $(LDFLAGS) -o $@ $^
```

The corresponding `advanced_example1.c` source file consists of a simple printing of a message, whose content depends on the preprocessor variable `MESSAGE`:

```
#include <stdio.h>

int main(){
#ifdef MESSAGE
    char *message = "SUCCESS";
#else
    char *message = "FAILURE";
#endif
    printf("Setting of preprocessor variable: %s\n", message);
    return 0;
}
```

The purpose of the regression test in this case is to set the preprocessor variable `MESSAGE` via `CPPFLAGS` and then check the standard output for the message `SUCCESS`, which indicates that the preprocessor flag has been passed and processed correctly by the Makefile.

The contents of this regression test are the following (tutorial/advanced/advanced_example1.py):

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class MakefileTest(rfm.RegressionTest):
    def __init__(self):
        self.descr = ('ReFrame tutorial demonstrating the use of Makefiles '
                     'and compile options')
        self.valid_systems = ['*']
        self.valid_prog_environs = ['*']
        self.executable = './advanced_example1'
        self.build_system = 'Make'
        self.build_system.cppflags = ['-DMESSAGE']
        self.sanity_patterns = sn.assert_found('SUCCESS', self.stdout)
        self.maintainers = ['put-your-name-here']
        self.tags = {'tutorial'}
```

The important bit here is how we set up the build system for this test:

```
self.build_system = 'Make'
self.build_system.cppflags = ['-DMESSAGE']
```

First, we set the build system to `Make` and then set the preprocessor flags for the compilation. ReFrame will invoke `make` as follows:

```
make -j 1 CC='cc' CXX='CC' FC='ftn' NVCC='nvcc' CPPFLAGS='-DMESSAGE'
```

The compiler variables (`CC`, `CXX` etc.) are set based on the corresponding values specified in the [configuration of the current environment](#). You may instruct the build system to ignore the default values from the environment by setting the following:

```
self.build_system.flags_from_environ = False
```

In this case, `make` will be invoked as follows:

```
make -j 1 CPPFLAGS='-DMESSAGE'
```

Notice that the `-j 1` option is always generated. You may change the maximum build concurrency as follows:

```
self.build_system.max_concurrency = 4
```

By setting `max_concurrency` to `None`, no limit for concurrent parallel jobs will be placed. This means that `make -j` will be used for building.

Finally, you may also customize the name of the Makefile. You can achieve that by setting the corresponding variable of the `Make` build system:

```
self.build_system.makefile = 'Makefile_custom'
```

More details on ReFrame’s build systems, you may find [here](#).

Retrieving the source code from a Git repository

It might be the case that a regression test needs to clone its source code from a remote repository. This can be achieved in two ways with ReFrame. One way is to set the `sourcedir` attribute to `None` and explicitly clone or checkout a repository using the `prebuild_cmd`:

```
self.sourcedir = None
self.prebuild_cmd = ['git clone https://github.com/me/myrepo .']
```

By setting `sourcedir` to `None`, you are telling ReFrame that you are going to provide the source files in the stage directory. The working directory of the `prebuild_cmd` and `postbuild_cmd` commands will be the stage directory of the test.

An alternative way to retrieve specifically a Git repository is to assign its URL directly to the `sourcedir` attribute:

```
self.sourcedir = 'https://github.com/me/myrepo'
```

ReFrame will attempt to clone this repository inside the stage directory by executing `git clone <repo> .` and will then proceed with the compilation as described above.

Note: ReFrame recognizes only URLs in the `sourcedir` attribute and requires passwordless access to the repository. This means that the SCP-style repository specification will not be accepted. You will have to specify it as URL using the `ssh://` protocol (see [Git documentation page](#)).

Add a configuration step before compiling the code

It is often the case that a configuration step is needed before compiling a code with `make`. To address this kind of projects, ReFrame aims to offer specific abstractions for “configure-make”-style build systems. It supports `CMake`-based projects through the `CMake` build system, as well as `Autotools`-based projects through the `Autotools` build system.

For other build systems, you can achieve the same effect using the `Make` build system and the `prebuild_cmd` for performing the configuration step. The following code snippet will configure a code with `./custom_configure` before invoking `make`:

```
self.prebuild_cmd = ['./custom_configure -with-mylib']
self.build_system = 'Make'
self.build_system.cppflags = ['-DHAVE_FOO']
self.build_system.flags_from_environ = False
```

The generated build script then will have the following lines:

```
./custom_configure -with-mylib
make -j 1 CPPFLAGS='-DHAVE_FOO'
```

3.5.2 Implementing a Run-Only Regression Test

There are cases when it is desirable to perform regression testing for an already built executable. The following test uses the `echo` Bash shell command to print a random integer between specific lower and upper bounds. Here is the full regression test (`tutorial/advanced/advanced_example2.py`):

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class ExampleRunOnlyTest(rfm.RunOnlyRegressionTest):
    def __init__(self):
        self.descr = ('ReFrame tutorial demonstrating the class'
                     'RunOnlyRegressionTest')
        self.valid_systems = ['*']
        self.valid_prog_environs = ['*']
        self.sourcedir = None

        lower = 90
        upper = 100
        self.executable = 'echo "Random: ${RANDOM%({1}+1-{})+{0}})".format(
            lower, upper)
        self.sanity_patterns = sn.assert_bounded(sn.extractsingle(
            r'Random: (?P<number>S+)', self.stdout, 'number', float),
            lower, upper)
        self.maintainers = ['put-your-name-here']
        self.tags = {'tutorial'}
```

There is nothing special for this test compared to those presented earlier except that it derives from the `RunOnlyRegressionTest` and that it does not contain any resources (`self.sourcedir = None`). Note that run-only regression tests may also have resources, as for instance a precompiled executable or some input data. The copying of these resources to the stage directory is performed at the beginning of the run phase. For standard regression tests, this happens at the beginning of the compilation phase, instead. Furthermore, in this particular test the `executable` consists only of standard Bash shell commands. For this reason, we can set `sourcedir` to `None` informing ReFrame that the test does not have any resources.

3.5.3 Implementing a Compile-Only Regression Test

ReFrame provides the option to write compile-only tests which consist only of a compilation phase without a specified executable. This kind of tests must derive from the `CompileOnlyRegressionTest` class provided by the framework. The following example (`tutorial/advanced/advanced_example3.py`) reuses the code of our first example in this section and checks that no warnings are issued by the compiler:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class ExampleCompileOnlyTest(rfm.CompileOnlyRegressionTest):
```

(continues on next page)

(continued from previous page)

```

def __init__(self):
    self.descr = ('ReFrame tutorial demonstrating the class'
                 'CompileOnlyRegressionTest')
    self.valid_systems = ['*']
    self.valid_prog_environs = ['*']
    self.sanity_patterns = sn.assert_not_found('warning', self.stderr)
    self.maintainers = ['put-your-name-here']
    self.tags = {'tutorial'}

```

The important thing to note here is that the standard output and standard error of the tests, accessible through the `stdout` and `stderr` attributes, are now the corresponding those of the compilation command. So sanity checking can be done in exactly the same way as with a normal test.

3.5.4 Leveraging Environment Variables

We have already demonstrated in the [tutorial](#) that ReFrame allows you to load the required modules for regression tests and also set any needed environment variables. When setting environment variables for your test through the `variables` attribute, you can assign them values of other, already defined, environment variables using the standard notation `$OTHER_VARIABLE` or `${OTHER_VARIABLE}`. The following regression test (tutorial/advanced/advanced_example4.py) sets the `CUDA_HOME` environment variable to the value of the `CUDAToolkit_HOME` and then compiles and runs a simple program:

```

import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class EnvironmentVariableTest(rfm.RegressionTest):
    def __init__(self):
        self.descr = ('ReFrame tutorial demonstrating the use'
                     'of environment variables provided by loaded modules')
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environs = ['*']
        self.modules = ['cudatoolkit']
        self.variables = {'CUDA_HOME': '$CUDAToolkit_HOME'}
        self.executable = './advanced_example4'
        self.build_system = 'Make'
        self.build_system.makefile = 'Makefile_example4'
        self.sanity_patterns = sn.assert_found(r'SUCCESS', self.stdout)
        self.maintainers = ['put-your-name-here']
        self.tags = {'tutorial'}

```

Before discussing this test in more detail, let's first have a look in the source code and the Makefile of this example:

```

#include <stdio.h>
#include <stdlib.h>
#include <string.h>

#ifdef CUDA_HOME
#   define CUDA_HOME ""
#endif

int main() {
    char *cuda_home_compile = CUDA_HOME;
    char *cuda_home_runtime = getenv("CUDA_HOME");

```

(continues on next page)

(continued from previous page)

```

if (cuda_home_runtime &&
    strlen(cuda_home_runtime, 256) &&
    strlen(cuda_home_compile, 256) &&
    !strcmp(cuda_home_compile, cuda_home_runtime, 256)) {
    printf("SUCCESS\n");
} else {
    printf("FAILURE\n");
    printf("Compiled with CUDA_HOME=%s, ran with CUDA_HOME=%s\n",
          cuda_home_compile,
          cuda_home_runtime ? cuda_home_runtime : "<null>");
}

return 0;
}

```

This program is pretty basic, but enough to demonstrate the use of environment variables from ReFrame. It simply compares the value of the `CUDA_HOME` macro with the value of the environment variable `CUDA_HOME` at runtime, printing `SUCCESS` if they are not empty and match. The Makefile for this example compiles this source by simply setting `CUDA_HOME` to the value of the `CUDA_HOME` environment variable:

```

EXECUTABLE := advanced_example4

CPPFLAGS = -DCUDA_HOME=\"$(CUDA_HOME)\"

.SUFFIXES: .o .c

OBJS := advanced_example4.o

$(EXECUTABLE): $(OBJS)
    $(CC) $(CFLAGS) $(LDFLAGS) -o $@ $^

$(OBJS): advanced_example4.c
    $(CC) $(CPPFLAGS) $(CFLAGS) -c $(LDFLAGS) -o $@ $^

clean:
    /bin/rm -f $(OBJS) $(EXECUTABLE)

```

Coming back now to the ReFrame regression test, the `CUDATOOLKIT_HOME` environment variable is defined by the `cuda` module. If you try to run the test, you will see that it will succeed, meaning that the `CUDA_HOME` variable was set correctly both during the compilation and the runtime.

When ReFrame sets up a test, it first loads its required modules and then sets the required environment variables expanding their values. This has the result that `CUDA_HOME` takes the correct value in our example at the compilation time.

At runtime, ReFrame will generate the following instructions in the shell script associated with this test:

```

module load cudatoolkit
export CUDA_HOME=$CUDATOOLKIT_HOME

```

This ensures that the environment of the test is also set correctly at runtime.

Finally, as already mentioned *previously*, since the name of the makefile is not one of the standard ones, it must be set explicitly in the build system:

```

self.build_system.makefile = 'Makefile_example4'

```


3.5.5 Setting a Time Limit for Regression Tests

ReFrame gives you the option to limit the execution time of regression tests. The following example (`tutorial/advanced/advanced_example5.py`) demonstrates how you can achieve this by limiting the execution time of a test that tries to sleep 100 seconds:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class TimeLimitTest(rfm.RunOnlyRegressionTest):
    def __init__(self):
        self.descr = ('ReFrame tutorial demonstrating the use'
                     'of a user-defined time limit')
        self.valid_systems = ['daint:gpu', 'daint:mc']
        self.valid_prog_environs = ['*']
        self.time_limit = (0, 1, 0)
        self.executable = 'sleep'
        self.executable_opts = ['100']
        self.sanity_patterns = sn.assert_found(
            r'CANCELLED.*DUE TO TIME LIMIT', self.stderr)
        self.maintainers = ['put-your-name-here']
        self.tags = {'tutorial'}
```

The important bit here is the following line that sets the time limit for the test to one minute:

```
self.time_limit = (0, 1, 0)
```

The `time_limit` attribute is a three-tuple in the form (HOURS, MINUTES, SECONDS). Time limits are implemented for all the scheduler backends.

The sanity condition for this test verifies that associated job has been canceled due to the time limit (note that this message is SLURM-specific).

```
self.sanity_patterns = sn.assert_found(
    r'CANCELLED.*DUE TO TIME LIMIT', self.stderr)
```

3.5.6 Applying a sanity function iteratively

It is often the case that a common sanity pattern has to be applied many times. In this example we will demonstrate how the above situation can be easily tackled using the `sanity` functions offered by ReFrame. Specifically, we would like to execute the following shell script and check that its output is correct:

```
#!/usr/bin/env bash

if [ -z $LOWER ]; then
    export LOWER=90
fi

if [ -z $UPPER ]; then
    export UPPER=100
fi

for i in {1..100}; do
```

(continues on next page)

```
echo Random: $((RANDOM%($UPPER+1-$LOWER)+$LOWER))
done
```

The above script simply prints 100 random integers between the limits given by the variables `LOWER` and `UPPER`. In the corresponding regression test we want to check that all the random numbers printed lie between 90 and 100 ensuring that the script executed correctly. Hence, a common sanity check has to be applied to all the printed random numbers. In ReFrame this can be achieved by the use of `map` sanity function accepting a function and an iterable as arguments. Through `map` the given function will be applied to all the members of the iterable object. Note that since `map` is a sanity function, its execution will be deferred. The contents of the ReFrame regression test contained in `advanced_example6.py` are the following:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class DeferredIterationTest(rfm.RunOnlyRegressionTest):
    def __init__(self):
        self.descr = ('ReFrame tutorial demonstrating the use of deferred '
                     'iteration via the `map` sanity function.')
        self.valid_systems = ['*']
        self.valid_prog_environs = ['*']
        self.executable = './random_numbers.sh'
        numbers = sn.extractall(
            r'Random: (?P<number>\S+)', self.stdout, 'number', float)
        self.sanity_patterns = sn.and_(
            sn.assert_eq(sn.count(numbers), 100),
            sn.all(sn.map(lambda x: sn.assert_bounded(x, 90, 100), numbers)))
        self.maintainers = ['put-your-name-here']
        self.tags = {'tutorial'}
```

First the random numbers are extracted through the `extractall` function as follows:

```
numbers = sn.extractall(
    r'Random: (?P<number>\S+)', self.stdout, 'number', float)
```

The `numbers` variable is a deferred iterable, which upon evaluation will return all the extracted numbers. In order to check that the extracted numbers lie within the specified limits, we make use of the `map` sanity function, which will apply the `assert_bounded` to all the elements of `numbers`. Additionally, our requirement is that all the numbers satisfy the above constraint and we therefore use `all`.

There is still a small complication that needs to be addressed. The `all` function returns `True` for empty iterables, which is not what we want. So we must ensure that all the numbers are extracted as well. To achieve this, we make use of `count` to get the number of elements contained in `numbers` combined with `assert_eq` to check that the number is indeed 100. Finally, both of the above conditions have to be satisfied for the program execution to be considered successful, hence the use of the `and_` function. Note that the `and` operator is not deferrable and will trigger the evaluation of any deferrable argument passed to it.

The full syntax for the `sanity_patterns` is the following:

```
self.sanity_patterns = sn.and_(
    sn.assert_eq(sn.count(numbers), 100),
    sn.all(sn.map(lambda x: sn.assert_bounded(x, 90, 100), numbers)))
```

3.5.7 Customizing the Generated Job Script

It is often the case that you must run some commands before and/or after the parallel launch of your executable. This can be easily achieved by using the `pre_run` and `post_run` attributes of `RegressionTest`.

The following example is a slightly modified version of the previous one. The lower and upper limits for the random numbers are now set inside a helper shell script in `scripts/limits.sh` and we want also to print the word `FINISHED` after our executable has finished. In order to achieve this, we need to source the helper script just before launching the executable and `echo` the desired message just after it finishes. Here is the test file:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class PrerunDemoTest(rfm.RunOnlyRegressionTest):
    def __init__(self):
        self.descr = ('ReFrame tutorial demonstrating the use of '
                     'pre- and post-run commands')
        self.valid_systems = ['*']
        self.valid_prog_environs = ['*']
        self.pre_run = ['source scripts/limits.sh']
        self.post_run = ['echo FINISHED']
        self.executable = './random_numbers.sh'
        numbers = sn.extractall(
            r'Random: (?P<number>\S+)', self.stdout, 'number', float)
        self.sanity_patterns = sn.all([
            sn.assert_eq(sn.count(numbers), 100),
            sn.all(sn.map(lambda x: sn.assert_bounded(x, 50, 80), numbers)),
            sn.assert_found('FINISHED', self.stdout)
        ])
        self.maintainers = ['put-your-name-here']
        self.tags = {'tutorial'}
```

Notice the use of the `pre_run` and `post_run` attributes. These are list of shell commands that are emitted verbatim in the job script. The generated job script for this example is the following:

```
#!/bin/bash -l
#SBATCH --job-name="prerun_demo_check_daint_gpu_PrgEnv-gnu"
#SBATCH --time=0:10:0
#SBATCH --ntasks=1
#SBATCH --output=prerun_demo_check.out
#SBATCH --error=prerun_demo_check.err
#SBATCH --constraint=gpu
module load daint-gpu
module unload PrgEnv-cray
module load PrgEnv-gnu
source scripts/limits.sh
srun ./random_numbers.sh
echo FINISHED
```

ReFrame generates the job shell script using the following pattern:

```
#!/bin/bash -l
{job_scheduler_preamble}
{test_environment}
{pre_run}
```

(continues on next page)

(continued from previous page)

```
{parallel_launcher} {executable} {executable_opts}
{post_run}
```

The `job_scheduler_preamble` contains the directives that control the job allocation. The `test_environment` are the necessary commands for setting up the environment of the test. This is the place where the modules and environment variables specified in `modules` and `variables` attributes are emitted. Then the commands specified in `pre_run` follow, while those specified in the `post_run` come after the launch of the parallel job. The parallel launch itself consists of three parts:

1. The parallel launcher program (e.g., `srun`, `mpirun` etc.) with its options,
2. the regression test executable as specified in the `executable` attribute and
3. the options to be passed to the executable as specified in the `executable_opts` attribute.

A key thing to note about the generated job script is that ReFrame submits it from the stage directory of the test, so that all relative paths are resolved against it.

3.5.8 Working with parameterized tests

New in version 2.13.

We have seen already in the [basic tutorial](#) how we could better organize the tests so as to avoid code duplication by using test class hierarchies. An alternative technique, which could also be used in parallel with the class hierarchies, is to use *parameterized tests*. The following is a test that takes a `variant` parameter, which controls which variant of the code will be used. Depending on that value, the test is set up differently:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.parameterized_test(['MPI'], ['OpenMP'])
class MatrixVectorTest(rfm.RegressionTest):
    def __init__(self, variant):
        self.descr = 'Matrix-vector multiplication test (%s)' % variant
        self.valid_systems = ['daint:gpu', 'daint:mc']
        self.valid_prog_environ = ['PrgEnv-cray', 'PrgEnv-gnu',
                                   'PrgEnv-intel', 'PrgEnv-pgi']

        self.build_system = 'SingleSource'
        self.prgenv_flags = {
            'PrgEnv-cray': ['-homp'],
            'PrgEnv-gnu': ['-fopenmp'],
            'PrgEnv-intel': ['-openmp'],
            'PrgEnv-pgi': ['-mp']
        }

        if variant == 'MPI':
            self.num_tasks = 8
            self.num_tasks_per_node = 2
            self.num_cpus_per_task = 4
            self.sourcepath = 'example_matrix_vector_multiplication_mpi_openmp.c'
        elif variant == 'OpenMP':
            self.sourcepath = 'example_matrix_vector_multiplication_openmp.c'
            self.num_cpus_per_task = 4

        self.variables = {
```

(continues on next page)

(continued from previous page)

```

        'OMP_NUM_THREADS': str(self.num_cpus_per_task)
    }
    matrix_dim = 1024
    iterations = 100
    self.executable_opts = [str(matrix_dim), str(iterations)]

    expected_norm = matrix_dim
    found_norm = sn.extractsingle(
        r'The L2 norm of the resulting vector is:\s+(?P<norm>\S+)',
        self.stdout, 'norm', float)
    self.sanity_patterns = sn.all([
        sn.assert_found(
            r'time for single matrix vector multiplication', self.stdout),
        sn.assert_lt(sn.abs(expected_norm - found_norm), 1.0e-6)
    ])
    self.maintainers = ['you-can-type-your-email-here']
    self.tags = {'tutorial'}

@rfm.run_before('compile')
def setflags(self):
    if self.prgenv_flags is not None:
        env = self.current_envIRON.name
        self.build_system.cflags = self.prgenv_flags[env]

```

If you have already gone through the [tutorial](#), this test can be easily understood. The new bit here is the `@parameterized_test` decorator of the `MatrixVectorTest` class. This decorator takes an arbitrary number of arguments, which are either of a sequence type (i.e., list, tuple etc.) or of a mapping type (i.e., dictionary). Each of the decorator's arguments corresponds to the constructor arguments of the decorated test that will be used to instantiate it. In the example shown, the test will be instantiated twice, once passing `variant` as `MPI` and a second time with `variant` passed as `OpenMP`. The framework will try to generate unique names for the generated tests by stringifying the arguments passed to the test's constructor:

```

Command line: ./bin/reframe -C tutorial/config/settings.py -c tutorial/advanced/
→advanced_example8.py -l
Reframe version: 2.15-dev1
Launched by user: XXX
Launched on host: daint101
Reframe paths
=====
    Check prefix      :
    Check search path : 'tutorial/advanced/advanced_example8.py'
    Stage dir prefix  : current/working/dir/reframe/stage/
    Output dir prefix : current/working/dir/reframe/output/
    Logging dir       : current/working/dir/reframe/logs
List of matched checks
=====
* MatrixVectorTest_MPI (Matrix-vector multiplication test (MPI))
* MatrixVectorTest_OpenMP (Matrix-vector multiplication test (OpenMP))
Found 2 check(s).

```

There are a couple of different ways that we could have used the `@parameterized_test` decorator. One is to use dictionaries for specifying the instantiations of our test class. The dictionaries will be converted to keyword arguments and passed to the constructor of the test class:

```
@rfm.parameterized_test({'variant': 'MPI'}, {'variant': 'OpenMP'})
```

Another way, which is quite useful if you want to generate lots of different tests at the same time, is to use either `list`

comprehensions or generator expressions for specifying the different test instantiations:

```
@rfm.parameterized_test(*([variant] for variant in ['MPI', 'OpenMP']))
```

Note: In versions of the framework prior to 2.13, this could be achieved by explicitly instantiating your tests inside the `_get_checks()` method.

Tip: Combining parameterized tests and test class hierarchies can offer you a very flexible way for generating multiple related tests at once keeping at the same time the maintenance cost low. We use this technique extensively in our tests.

3.5.9 Flexible Regression Tests

New in version 2.15.

ReFrame can automatically set the number of tasks of a particular test, if its `num_tasks` attribute is set to ≤ 0 . In ReFrame's terminology, such tests are called *flexible*. Negative values indicate the minimum number of tasks that is acceptable for this test (a value of -4 indicates a minimum acceptable number of 4 tasks). A zero value indicates the default minimum number of tasks which is equal to `num_tasks_per_node`.

By default, ReFrame will spawn such a test on all the idle nodes of the current system partition, but this behavior can be adjusted from the command-line. Flexible tests are very useful for diagnostics tests, e.g., tests for checking the health of a whole set nodes. In this example, we demonstrate this feature through a simple test that runs `hostname`. The test will verify that all the nodes print the expected host name:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class HostnameCheck(rfm.RunOnlyRegressionTest):
    def __init__(self):
        self.valid_systems = ['daint:gpu', 'daint:mc']
        self.valid_prog_environ = ['PrgEnv-cray']
        self.executable = 'hostname'
        self.sourcesdir = None
        self.num_tasks = 0
        self.num_tasks_per_node = 1
        self.sanity_patterns = sn.assert_eq(
            self.num_tasks_assigned,
            sn.count(sn.findall(r'nid\d+', self.stdout))
        )
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}

    @property
    @sn.sanity_function
    def num_tasks_assigned(self):
        return self.job.num_tasks
```

The first thing to notice in this test is that `num_tasks` is set to 0. This is a requirement for flexible tests:

```
self.num_tasks = 0
```

The sanity function of this test simply counts the host names and verifies that they are as many as expected:

```
self.sanity_patterns = sn.assert_eq(
    self.num_tasks_assigned,
    sn.count(sn.findall(r'nid\d+', self.stdout))
)
```

Notice, however, that the sanity check does not use `num_tasks` for verification, but rather a different, custom attribute, the `num_tasks_assigned`. This happens for two reasons:

- At the time the sanity check expression is created, `num_tasks` is 0. So the actual number of tasks assigned must be a deferred expression as well.
- When ReFrame will determine and set the number of tasks of the test, it will not set the `num_tasks` attribute of the `RegressionTest`. It will only set the corresponding attribute of the associated job instance.

Here is how the new deferred attribute is defined:

```
@property
@sn.sanity_function
def num_tasks_assigned(self):
    return self.job.num_tasks
```

The behavior of the flexible task allocation is controlled by the `--flex-alloc-nodes` command line option. See the corresponding [section](#) for more information.

3.5.10 Testing containerized applications

New in version 2.20.

ReFrame can be used also to test applications that run inside a container. A container-based test can be written as `RunOnlyRegressionTest` that sets the `container_platform`. The following example shows a simple test that runs some basic commands inside an Ubuntu 18.04 container and checks that the test has indeed run inside the container and that the stage directory was correctly mounted:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.required_version('>=2.20-dev2')
@rfm.simple_test
class Example10Test(rfm.RunOnlyRegressionTest):
    def __init__(self):
        self.descr = 'Run commands inside a container'
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environs = ['PrgEnv-cray']
        self.container_platform = 'Singularity'
        self.container_platform.image = 'docker://ubuntu:18.04'
        self.container_platform.commands = [
            'pwd', 'ls', 'cat /etc/os-release'
        ]
        self.container_platform.workdir = '/workdir'
        self.sanity_patterns = sn.all([
            sn.assert_found(r'^' + self.container_platform.workdir,
                           self.stdout),
            sn.assert_found(r'^advanced_example1.c', self.stdout),
            sn.assert_found(r'18.04.\d+ LTS \(Bionic Beaver\) ', self.stdout),
```

(continues on next page)

(continued from previous page)

```

    ])
    self.maintainers = ['put-your-name-here']
    self.tags = {'tutorial'}

```

A container-based test in ReFrame requires that the `container_platform` is set:

```
self.container_platform = 'Singularity'
```

This attribute accepts a string that corresponds to the name of the platform and it instantiates the appropriate `ContainerPlatform` object behind the scenes. In this case, the test will be using `Singularity` as a container platform. If such a platform is not configured for the current system, the test will fail. For a complete list of supported container platforms, the user is referred to the [configuration documentation](#).

As soon as the container platform to be used is defined, you need to specify the container image to use and the commands to run inside the container:

```

self.container_platform.image = 'docker://ubuntu:18.04'
self.container_platform.commands = [
    'pwd', 'ls', 'cat /etc/os-release'
]

```

These two attributes are mandatory for container-based check. The `image` attribute specifies the name of an image from a registry, whereas the `commands` attribute provides the list of commands to be run inside the container. It is important to note that the `executable` and `executable_opts` attributes of the `RegressionTest` are ignored in case of container-based tests.

In the above example, ReFrame will run the container as follows:

```

singularity exec -B"/path/to/test/stagedir:/rfm_workdir" docker://ubuntu:18.04 bash -
↪c 'cd rfm_workdir; pwd; ls; cat /etc/os-release'

```

By default ReFrame will mount the stage directory of the test under `/rfm_workdir` inside the container and it will always prepend a `cd` command to that directory. The user commands then are then run from that directory one after the other. Once the commands are executed, the container is stopped and ReFrame goes on with the sanity and/or performance checks.

Users may also change the default mount point of the stage directory by using `workdir` attribute:

```
self.container_platform.workdir = '/workdir'
```

Besides the stage directory, additional mount points can be specified through the `mount_points` attribute:

```

self.container_platform.mount_points = [('/path/to/host/dir1', '/path/to/container/
↪mount_point1'),
                                       ('/path/to/host/dir2', '/path/to/container/
↪mount_point2')]

```

For a complete list of the available attributes of a specific container platform, the reader is referred to [reference guide](#).

3.5.11 Using dependencies in your tests

New in version 2.21.

A ReFrame test may define dependencies to other tests. An example scenario is to test different runtime configurations of a benchmark that you need to compile, or run a scaling analysis of a code. In such cases, you don't want to rebuild your test for each runtime configuration. You could have a build test, which all runtime tests would depend on. This

is the approach we take with the following example, that fetches, builds and runs several OSU benchmarks. We first create a basic compile-only test, that fetches the benchmarks and builds them for the different programming environments:

```
@rfm.simple_test
class OSUBuildTest(rfm.CompileOnlyRegressionTest):
    def __init__(self):
        self.descr = 'OSU benchmarks build test'
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environs = ['PrgEnv-gnu', 'PrgEnv-pgi', 'PrgEnv-intel']
        self.sourcedir = None
        self.prebuild_cmd = [
            'wget http://mvapich.cse.ohio-state.edu/download/mvapich/osu-micro-
            ↪benchmarks-5.6.2.tar.gz',
            'tar xzf osu-micro-benchmarks-5.6.2.tar.gz',
            'cd osu-micro-benchmarks-5.6.2'
        ]
        self.build_system = 'Autotools'
        self.build_system.max_concurrency = 8
        self.sanity_patterns = sn.assert_not_found('error', self.stderr)
```

There is nothing particular to that test, except perhaps that you can set `sourcedir` to `None` even for a test that needs to compile something. In such a case, you should at least provide the commands that fetch the code inside the `prebuild_cmd` attribute.

For the next test we need to use the OSU benchmark binaries that we just built, so as to run the MPI ping-pong benchmark. Here is the relevant part:

```
class OSUBenchmarkTestBase(rfm.RunOnlyRegressionTest):
    '''Base class of OSU benchmarks runtime tests'''

    def __init__(self):
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environs = ['PrgEnv-gnu', 'PrgEnv-pgi', 'PrgEnv-intel']
        self.sourcedir = None
        self.num_tasks = 2
        self.num_tasks_per_node = 1
        self.sanity_patterns = sn.assert_found(r'^8', self.stdout)

@rfm.simple_test
class OSULatencyTest(OSUBenchmarkTestBase):
    def __init__(self):
        super().__init__()
        self.descr = 'OSU latency test'
        self.perf_patterns = {
            'latency': sn.extractsingle(r'^8\s+(\S+)', self.stdout, 1, float)
        }
        self.depends_on('OSUBuildTest')
        self.reference = {
            '*': {'latency': (0, None, None, 'us')}
        }

@rfm.require_deps
def set_executable(self, OSUBuildTest):
    self.executable = os.path.join(
        OSUBuildTest().stagedir,
        'osu-micro-benchmarks-5.6.2', 'mpi', 'pt2pt', 'osu_latency'
```

(continues on next page)

(continued from previous page)

```
)
self.executable_opts = ['-x', '100', '-i', '1000']
```

First, since we will have multiple similar benchmarks, we move all the common functionality to the `OSUBenchmarkTestBase` base class. Again nothing new here; we are going to use two nodes for the benchmark and we set `sourcedir` to `None`, since none of the benchmark tests will use any additional resources. The new part comes in with the `OSULatencyTest` test in the following line:

```
self.depends_on('OSUBuildTest')
```

Here we tell ReFrame that this test depends on a test named `OSUBuildTest`. This test may or may not be defined in the same test file; all ReFrame needs is the test name. By default, the `depends_on()` function will create dependencies between the individual test cases of the `OSULatencyTest` and the `OSUBuildTest`, such that the `OSULatencyTest` using `PrgEnv-gnu` will depend on the outcome of the `OSUBuildTest` using `PrgEnv-gnu`, but not on the outcome of the `OSUBuildTest` using `PrgEnv-intel`. This behaviour can be changed, but we will return to this later. You can create arbitrary test dependency graphs, but they need to be acyclic. If ReFrame detects cyclic dependencies, it will refuse to execute the set of tests and will issue an error pointing out the cycle.

A ReFrame test with dependencies will execute, i.e., enter its *setup* stage, only after *all* of its dependencies have succeeded. If any of its dependencies fails, the current test will be marked as failure as well.

The next step for the `OSULatencyTest` is to set its executable to point to the binary produced by the `OSUBuildTest`. This is achieved with the following specially decorated function:

```
@rfm.require_deps
def set_executable(self, OSUBuildTest):
    self.executable = os.path.join(
        OSUBuildTest().stagedir,
        'osu-micro-benchmarks-5.6.2', 'mpi', 'pt2pt', 'osu_latency'
    )
    self.executable_opts = ['-x', '100', '-i', '1000']
```

The `@require_deps` decorator will bind the arguments passed to the decorated function to the result of the dependency that each argument names. In this case, it binds the `OSUBuildTest` function argument to the result of a dependency named `OSUBuildTest`. In order for the binding to work correctly the function arguments must be named after the target dependencies. However, referring to a dependency only by the test's name is not enough, since a test might be associated with multiple programming environments. For this reason, a dependency argument is actually bound to a function that accepts as argument the name of a target programming environment. If no arguments are passed to that function, as in this example, the current programming environment is implied, such that `OSUBuildTest()` is equivalent to `OSUBuildTest(self.current_env.name)`. This call returns the actual test case of the dependency that has been executed. This allows you to access any attribute from the target test, as we do in this example by accessing the target test's stage directory, which we use to construct the path of the executable. This concludes the presentation of the `OSULatencyTest` test. The `OSUBandwidthTest` is completely analogous.

The `OSUAllreduceTest` shown below is similar to the other two, except that it is parameterized. It is essentially a scalability test that is running the `osu_allreduce` executable created by the `OSUBuildTest` for 2, 4, 8 and 16 nodes.

```
@rfm.parameterized_test(*([1 << i] for i in range(1, 5)))
class OSUAllreduceTest(OSUBenchmarkTestBase):
    def __init__(self, num_tasks):
        super().__init__()
        self.descr = 'OSU Allreduce test'
```

(continues on next page)

(continued from previous page)

```

self.perf_patterns = {
    'latency': sn.extractsingle(r'^8\s+(\S+)', self.stdout, 1, float)
}
self.depends_on('OSUBuildTest')
self.reference = {
    '*': {'latency': (0, None, None, 'us')}
}
self.num_tasks = num_tasks

@rfm.require_deps
def set_executable(self, OSUBuildTest):
    self.executable = os.path.join(
        OSUBuildTest().stagedir,
        'osu-micro-benchmarks-5.6.2', 'mpi', 'collective', 'osu_allreduce'
    )
    self.executable_opts = ['-m', '8', '-x', '1000', '-i', '20000']

```

The full set of OSU example tests is shown below:

```

import os

import reframe as rfm
import reframe.utility.sanity as sn

class OSUBenchmarkTestBase(rfm.RunOnlyRegressionTest):
    '''Base class of OSU benchmarks runtime tests'''

    def __init__(self):
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environs = ['PrgEnv-gnu', 'PrgEnv-pgi', 'PrgEnv-intel']
        self.sourcesdir = None
        self.num_tasks = 2
        self.num_tasks_per_node = 1
        self.sanity_patterns = sn.assert_found(r'^8', self.stdout)

@rfm.simple_test
class OSULatencyTest(OSUBenchmarkTestBase):
    def __init__(self):
        super().__init__()
        self.descr = 'OSU latency test'
        self.perf_patterns = {
            'latency': sn.extractsingle(r'^8\s+(\S+)', self.stdout, 1, float)
        }
        self.depends_on('OSUBuildTest')
        self.reference = {
            '*': {'latency': (0, None, None, 'us')}
        }

    @rfm.require_deps
    def set_executable(self, OSUBuildTest):
        self.executable = os.path.join(
            OSUBuildTest().stagedir,
            'osu-micro-benchmarks-5.6.2', 'mpi', 'pt2pt', 'osu_latency'
        )
        self.executable_opts = ['-x', '100', '-i', '1000']

```

(continues on next page)

```

@rfm.simple_test
class OSUBandwidthTest (OSUBenchmarkTestBase):
    def __init__(self):
        super().__init__()
        self.descr = 'OSU bandwidth test'
        self.perf_patterns = {
            'bandwidth': sn.extractsingle(r'^4194304\s+(\S+)',
                                         self.stdout, 1, float)
        }
        self.depends_on('OSUBuildTest')
        self.reference = {
            '*': {'bandwidth': (0, None, None, 'MB/s')}
        }

    @rfm.require_deps
    def set_executable(self, OSUBuildTest):
        self.executable = os.path.join(
            OSUBuildTest().stagedir,
            'osu-micro-benchmarks-5.6.2', 'mpi', 'pt2pt', 'osu_bw'
        )
        self.executable_opts = ['-x', '100', '-i', '1000']

@rfm.parameterized_test(*( [1 << i] for i in range(1, 5)))
class OSUAllreduceTest (OSUBenchmarkTestBase):
    def __init__(self, num_tasks):
        super().__init__()
        self.descr = 'OSU Allreduce test'
        self.perf_patterns = {
            'latency': sn.extractsingle(r'^8\s+(\S+)', self.stdout, 1, float)
        }
        self.depends_on('OSUBuildTest')
        self.reference = {
            '*': {'latency': (0, None, None, 'us')}
        }
        self.num_tasks = num_tasks

    @rfm.require_deps
    def set_executable(self, OSUBuildTest):
        self.executable = os.path.join(
            OSUBuildTest().stagedir,
            'osu-micro-benchmarks-5.6.2', 'mpi', 'collective', 'osu_allreduce'
        )
        self.executable_opts = ['-m', '8', '-x', '1000', '-i', '20000']

@rfm.simple_test
class OSUBuildTest (rfm.CompileOnlyRegressionTest):
    def __init__(self):
        self.descr = 'OSU benchmarks build test'
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environ = ['PrgEnv-gnu', 'PrgEnv-pgi', 'PrgEnv-intel']
        self.sourcedir = None
        self.prebuild_cmd = [
            'wget http://mvapich.cse.ohio-state.edu/download/mvapich/osu-micro-
↪benchmarks-5.6.2.tar.gz',

```

(continues on next page)

(continued from previous page)

```

        'tar xzf osu-micro-benchmarks-5.6.2.tar.gz',
        'cd osu-micro-benchmarks-5.6.2'
    ]
    self.build_system = 'Autotools'
    self.build_system.max_concurrency = 8
    self.sanity_patterns = sn.assert_not_found('error', self.stderr)

```

Notice that the order dependencies are defined in a test file is irrelevant. In this case, we define OSUBuildTest at the end. ReFrame will make sure to properly sort the tests and execute them.

Here is the output when running the OSU tests with the asynchronous execution policy:

```

[=====] Running 7 check(s)
[=====] Started on Tue Dec 10 00:15:53 2019

[-----] started processing OSUBuildTest (OSU benchmarks build test)
[ RUN    ] OSUBuildTest on daint:gpu using PrgEnv-gnu
[ RUN    ] OSUBuildTest on daint:gpu using PrgEnv-intel
[ RUN    ] OSUBuildTest on daint:gpu using PrgEnv-pgi
[-----] finished processing OSUBuildTest (OSU benchmarks build test)

[-----] started processing OSULatencyTest (OSU latency test)
[ RUN    ] OSULatencyTest on daint:gpu using PrgEnv-gnu
[   DEP  ] OSULatencyTest on daint:gpu using PrgEnv-gnu
[ RUN    ] OSULatencyTest on daint:gpu using PrgEnv-intel
[   DEP  ] OSULatencyTest on daint:gpu using PrgEnv-intel
[ RUN    ] OSULatencyTest on daint:gpu using PrgEnv-pgi
[   DEP  ] OSULatencyTest on daint:gpu using PrgEnv-pgi
[-----] finished processing OSULatencyTest (OSU latency test)

[-----] started processing OSUBandwidthTest (OSU bandwidth test)
[ RUN    ] OSUBandwidthTest on daint:gpu using PrgEnv-gnu
[   DEP  ] OSUBandwidthTest on daint:gpu using PrgEnv-gnu
[ RUN    ] OSUBandwidthTest on daint:gpu using PrgEnv-intel
[   DEP  ] OSUBandwidthTest on daint:gpu using PrgEnv-intel
[ RUN    ] OSUBandwidthTest on daint:gpu using PrgEnv-pgi
[   DEP  ] OSUBandwidthTest on daint:gpu using PrgEnv-pgi
[-----] finished processing OSUBandwidthTest (OSU bandwidth test)

[-----] started processing OSUAllreduceTest_2 (OSU Allreduce test)
[ RUN    ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-gnu
[   DEP  ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-gnu
[ RUN    ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-intel
[   DEP  ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-intel
[ RUN    ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-pgi
[   DEP  ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-pgi
[-----] finished processing OSUAllreduceTest_2 (OSU Allreduce test)

[-----] started processing OSUAllreduceTest_4 (OSU Allreduce test)
[ RUN    ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-gnu
[   DEP  ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-gnu
[ RUN    ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-intel
[   DEP  ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-intel
[ RUN    ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-pgi
[   DEP  ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-pgi
[-----] finished processing OSUAllreduceTest_4 (OSU Allreduce test)

```

(continues on next page)

(continued from previous page)

```

[-----] started processing OSUAllreduceTest_8 (OSU Allreduce test)
[ RUN    ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-gnu
[   DEP  ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-gnu
[ RUN    ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-intel
[   DEP  ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-intel
[ RUN    ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-pgi
[   DEP  ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-pgi
[-----] finished processing OSUAllreduceTest_8 (OSU Allreduce test)

[-----] started processing OSUAllreduceTest_16 (OSU Allreduce test)
[ RUN    ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-gnu
[   DEP  ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-gnu
[ RUN    ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-intel
[   DEP  ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-intel
[ RUN    ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-pgi
[   DEP  ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-pgi
[-----] finished processing OSUAllreduceTest_16 (OSU Allreduce test)

[-----] waiting for spawned checks to finish
[   OK   ] OSUBuildTest on daint:gpu using PrgEnv-pgi
[   OK   ] OSUBuildTest on daint:gpu using PrgEnv-gnu
[   OK   ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-pgi
[   OK   ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-gnu
[   OK   ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-gnu
[   OK   ] OSUBuildTest on daint:gpu using PrgEnv-intel
[   OK   ] OSULatencyTest on daint:gpu using PrgEnv-gnu
[   OK   ] OSUBandwidthTest on daint:gpu using PrgEnv-gnu
[   OK   ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-gnu
[   OK   ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-pgi
[   OK   ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-pgi
[   OK   ] OSULatencyTest on daint:gpu using PrgEnv-intel
[   OK   ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-intel
[   OK   ] OSUAllreduceTest_16 on daint:gpu using PrgEnv-intel
[   OK   ] OSUBandwidthTest on daint:gpu using PrgEnv-pgi
[   OK   ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-pgi
[   OK   ] OSUAllreduceTest_8 on daint:gpu using PrgEnv-intel
[   OK   ] OSUAllreduceTest_4 on daint:gpu using PrgEnv-gnu
[   OK   ] OSULatencyTest on daint:gpu using PrgEnv-pgi
[   OK   ] OSUAllreduceTest_2 on daint:gpu using PrgEnv-intel
[   OK   ] OSUBandwidthTest on daint:gpu using PrgEnv-intel
[-----] all spawned checks have finished

[ PASSED ] Ran 21 test case(s) from 7 check(s) (0 failure(s))
[=====] Finished on Tue Dec 10 00:21:11 2019

```

Before starting running the tests, ReFrame topologically sorts them based on their dependencies and schedules them for running using the selected execution policy. With the serial execution policy, ReFrame simply executes the tests to completion as they “arrive”, since the tests are already topologically sorted. In the asynchronous execution policy, tests are spawned and not waited for. If a test’s dependencies have not yet completed, it will not start its execution and a DEP message will be printed to denote this.

Finally, ReFrame’s runtime takes care of properly cleaning up the resources of the tests respecting dependencies. Normally when an individual test finishes successfully, its stage directory is cleaned up. However, if other tests are depending on this one, this would be catastrophic, since most probably the dependent tests would need the outcome of this test. ReFrame fixes that by not cleaning up the stage directory of a test until all its dependent tests have finished successfully.

3.6 How Test Dependencies Work In ReFrame

Before going into details on how ReFrame treats test dependencies, it is important to understand how tests are actually treated and executed by the runtime. Normally, a ReFrame test will be tried for different programming environments and different partitions within the same ReFrame run. These are defined in the test's `__init__()` method, but it is not this original object that is being executed by the regression test pipeline. The following figure explains in more detail the process:

When ReFrame loads a test from the disk it unconditionally constructs it executing its `__init__()` method. The practical implication of this is that your test will be instantiated even if it will not run on the current system. After all the tests are loaded, they are filtered based on the current system and any other criteria (such as programming environment, test attributes etc.) specified by the user (see [Filtering of Regression Tests](#) for more details). After the tests are filtered, ReFrame creates the actual *test cases* to be run. A test case is essentially a tuple consisting of the test, the system partition and the programming environment to try. The test that goes into a test case is essentially a *clone* of the original test that was instantiated upon loading. This ensures that the test case's state is not shared and may not be reused in any case. Finally, the generated test cases are passed to a *runner* that is responsible for scheduling them for execution based on the selected execution policy.

Dependencies in ReFrame are defined at the test level using the `depends_on()` function, but are projected to the test cases space. We will see the rules of that projection in a while. The dependency graph construction and the subsequent dependency analysis happen also at the level of the test cases.

Let's assume that test T1 depends in T0. This can be expressed inside T1 using the `depends_on()` method:

```
@rfm.simple_test
class T1(rfm.RegressionTest):
    def __init__(self):
        ...
        self.depends_on('T0')
```

Conceptually, this dependency can be viewed at the test level as follows:

For most of the cases, this is sufficient to reason about test dependencies. In reality, as mentioned above, dependencies are handled at the level of test cases. Test cases on different partitions are always independent. If not specified differently, test cases using programming environments are also independent. This is the default behavior of the `depends_on()` function. The following image shows the actual test case dependencies assuming that both tests support the E0 and E1 programming environments (for simplicity, we have omitted the partitions, since tests are always independent in that dimension):

This means that test cases of T1 may start executing before all test cases of T0 have finished. You can impose a stricter dependency between tests, such that T1 does not start execution unless all test cases of T0 have finished. You can achieve this as follows:

```
@rfm.simple_test
class T1(rfm.RegressionTest):
    def __init__(self):
        ...
        self.depends_on('T0', how=rfm.DEPEND_FULLY)
```

This will create the following test case graph:

You may also create arbitrary dependencies between the test cases of different tests, like in the following example, where the dependencies cannot be represented in any of the other two ways:

These dependencies can be achieved as follows:

```
@rfm.simple_test
class T1(rfm.RegressionTest):
    def __init__(self):
        ...
        self.depends_on('T0', how=rfm.DEPEND_EXACT,
                        subdeps={'E0': ['E0', 'E1'], 'E1': ['E1']})
```

The `subdeps` argument defines the sub-dependencies between the test cases of T1 and T0 using an adjacency list representation.

3.6.1 Cyclic dependencies

Obviously, cyclic dependencies between test cases are not allowed. Cyclic dependencies between tests are not allowed either, even if the test case dependency graph is acyclic. For example, the following dependency set up is invalid:

The test case dependencies here, clearly, do not form a cycle, but the edge from (T0, E0) to (T1, E1) introduces a dependency from T0 to T1 forming a cycle at the test level. The reason we impose this restriction is that we wanted to keep the original processing of tests by ReFrame, where all the test cases of a test are processed before moving to the next one. Supporting this type of dependencies would require to change substantially ReFrame's output.

3.6.2 Dangling dependencies

In our discussion so far, T0 and T1 had the same valid programming environments. What happens if they do not? Assume, for example, that T0 and T1 are defined as follows:

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class T0(rfm.RegressionTest):
    def __init__(self):
        self.valid_systems = ['P0']
        self.valid_prog_environs = ['E0']
        ...

@rfm.simple_test
class T1(rfm.RegressionTest):
    def __init__(self):
        self.valid_systems = ['P0']
        self.valid_prog_environs = ['E0', 'E1']
        self.depends_on('T0')
        ...
```

As discussed previously, `depends_on()` will create one-to-one dependencies between the different programming environment test cases. So in this case it will try to create an edge from (T1, E1) to (T0, E1) as shown below:

This edge cannot be resolved since the target test case does not exist. ReFrame will complain and issue an error while trying to build the test dependency graph. The remedy to this is to use either `DEPEND_FULLY` or pass the exact dependencies with `DEPEND_EXACT` to `depends_on()`.

If T0 and T1 had their `valid_prog_environs` swapped, such that T0 supported E0 and E1 and T1 supported only E0, the default `depends_on()` mode would work fine. The (T0, E1) test case would simply have no dependent test cases.

3.6.3 Resolving dependencies

As shown in the [tutorial](#), test dependencies would be of limited usage if you were not able to use the results or information of the target tests. Let's reiterate over the `set_executable()` function of the `OSULatencyTest` that we presented previously:

```
@rfm.require_deps
def set_executable(self, OSUBuildTest):
    self.executable = os.path.join(
        OSUBuildTest().stagedir,
        'osu-micro-benchmarks-5.6.2', 'mpi', 'pt2pt', 'osu_latency'
    )
    self.executable_opts = ['-x', '100', '-i', '1000']
```

The `@require_deps` decorator does some magic – we will unravel this shortly – with the function arguments of the `set_executable()` function and binds them to the target test dependencies by their name. However, as discussed in this section, dependencies are defined at test case level, so the `OSUBuildTest` function argument is bound to a special function that allows you to retrieve an actual test case of the target dependency. This is why you need to “call” `OSUBuildTest` in order to retrieve the desired test case. When no arguments are passed, this will retrieve the test case corresponding to the current partition and the current programming environment. We could always retrieve the `PrgEnv-gnu` case by writing `OSUBuildTest('PrgEnv-gnu')`. If a dependency cannot be resolved, because it is invalid, a runtime error will be thrown with an appropriate message.

The low-level method for retrieving a dependency is the `getdep()` method of the `RegressionTest`. In fact, you can rewrite `set_executable()` function as follows:

```
@rfm.run_after('setup')
def set_executable(self):
    target = self.getdep('OSUBuildTest')
    self.executable = os.path.join(
        target.stagedir,
        'osu-micro-benchmarks-5.6.2', 'mpi', 'pt2pt', 'osu_latency'
    )
    self.executable_opts = ['-x', '100', '-i', '1000']
```

Now it's easier to understand what the `@require_deps` decorator does behind the scenes. It binds the function arguments to a partial realization of the `getdep()` function and attaches the decorated function as an after-setup hook. In fact, any `@require_deps`-decorated function will be invoked before any other after-setup hook.

3.7 Understanding the Mechanism of Sanity Functions

This section describes the mechanism behind the sanity functions that are used for the sanity and performance checking. Generally, writing a new sanity function is as straightforward as decorating a simple Python function with either

the `sanity_function` or the `@reframe.core.deferrable.deferrable` decorator. However, it is important to understand how and when a deferrable function is evaluated, especially if your function takes as arguments the results of other deferrable functions.

3.7.1 What Is a Deferrable Function?

A deferrable function is a function whose evaluation is deferred to a later point in time. You can define any function as deferrable by adding the `@sanity_function` or the `@deferrable` decorator before its definition. The example below demonstrates a simple scenario:

```
import reframe.utility.sanity as sn

@sn.sanity_function
def foo():
    print('hello')
```

If you try to call `foo()`, its code will not execute:

```
>>> foo()
<reframe.core.deferrable._DeferredExpression object at 0x2b70fff23550>
```

Instead, a special object is returned that represents the function whose execution is deferred. Notice the more general *deferred expression* name of this object. We shall see later on why this name is used.

In order to explicitly trigger the execution of `foo()`, you have to call *evaluate* on it:

```
>>> from reframe.utility.sanity import evaluate
>>> evaluate(foo())
hello
```

If the argument passed to *evaluate* is not a deferred expression, it will be simply returned as is.

Deferrable functions may also be combined as we do with normal functions. Let's extend our example with `foo()` accepting an argument and printing it:

```
import reframe.utility.sanity as sn

@sn.sanity_function
def foo(arg):
    print(arg)

@sn.sanity_function
def greetings():
    return 'hello'
```

If we now do `foo(greetings())`, again nothing will be evaluated:

```
>>> foo(greetings())
<reframe.core.deferrable._DeferredExpression object at 0x2b7100e9e978>
```

If we trigger the evaluation of `foo()` as before, we will get expected result:

```
>>> evaluate(foo(greetings()))
hello
```

Notice how the evaluation mechanism goes down the function call graph and returns the expected result. An alternative way to evaluate this expression would be the following:

```
>>> x = foo(greetings())
>>> x.evaluate()
hello
```

As you may have noticed, you can assign a deferred function to a variable and evaluate it later. You may also do `evaluate(x)`, which is equivalent to `x.evaluate()`.

To demonstrate more clearly how the deferred evaluation of a function works, let's consider the following `size3()` deferrable function that simply checks whether an `iterable` passed as argument has three elements inside it:

```
@sn.sanity_function
def size3(iterable):
    return len(iterable) == 3
```

Now let's assume the following example:

```
>>> l = [1, 2]
>>> x = size3(l)
>>> evaluate(x)
False
>>> l += [3]
>>> evaluate(x)
True
```

We first call `size3()` and store its result in `x`. As expected when we evaluate `x`, `False` is returned, since at the time of the evaluation our list has two elements. We later append an element to our list and reevaluate `x` and we get `True`, since at this point the list has three elements.

Note: Deferred functions and expressions may be stored and (re)evaluated at any later point in the program.

An important thing to point out here is that deferrable functions *capture* their arguments at the point they are called. If you change the binding of a variable name (either explicitly or implicitly by applying an operator to an immutable object), this change will not be reflected when you evaluate the deferred function. The function instead will operate on its captured arguments. We will demonstrate this by replacing the list in the above example with a tuple:

```
>>> l = (1, 2)
>>> x = size3(l)
>>> l += (3,)
>>> l
(1, 2, 3)
>>> evaluate(x)
False
```

Why this is happening? This is because tuples are immutable so when we are doing `l += (3,)` to append to our tuple, Python constructs a new tuple and rebinds `l` to the newly created tuple that has three elements. However, when we called our deferrable function, `l` was pointing to a different tuple object, and that was the actual tuple argument that our deferrable function has captured.

The following augmented example demonstrates this:

```
>>> l = (1, 2)
>>> x = size3(l)
>>> l += (3,)
>>> l
(1, 2, 3)
>>> evaluate(x)
```

(continues on next page)

(continued from previous page)

```
False
>>> l = (1, 2)
>>> id(l)
47764346657160
>>> x = size3(l)
>>> l += (3,)
>>> id(l)
47764330582232
>>> l
(1, 2, 3)
>>> evaluate(x)
False
```

Notice the different IDs of `l` before and after the `+=` operation. This is a key trait of deferrable functions and expressions that you should be aware of.

3.7.2 Deferred expressions

You might be still wondering why the internal name of a deferred function refers to the more general term deferred expression. Here is why:

```
>>> @sn.sanity_function
... def size(iterable):
...     return len(iterable)
...
>>> l = [1, 2]
>>> x = 2*(size(l) + 3)
>>> x
<reframe.core.deferrable._DeferredExpression object at 0x2b1288f4e940>
>>> evaluate(x)
10
```

As you can see, you can use the result of a deferred function inside arithmetic operations. The result will be another deferred expression that you can evaluate later. You can practically use any Python builtin operator or builtin function with a deferred expression and the result will be another deferred expression. This is quite a powerful mechanism, since with the standard syntax you can create arbitrary expressions that may be evaluated later in your program.

There are some exceptions to this rule, though. The logical `and`, `or` and `not` operators as well as the `in` operator cannot be deferred automatically. These operators try to take the truthy value of their arguments by calling `bool` on them. As we shall see later, applying the `bool` function on a deferred expression causes its immediate evaluation and returns the result. If you want to defer the execution of such operators, you should use the corresponding `and_`, `or_`, `not_` and `contains` functions in `reframe.utility.sanity`, which basically wrap the expression in a deferrable function.

In summary deferrable functions have the following characteristics:

- You can make any function deferrable by preceding it with the `@sanity_function` or the `@deferrable` decorator.
- When you call a deferrable function, its body is not executed but its arguments are *captured* and an object representing the deferred function is returned.
- You can execute the body of a deferrable function at any later point by calling `evaluate` on the deferred expression object that it has been returned by the call to the deferred function.
- Deferred functions can accept other deferred expressions as arguments and may also return a deferred expression.

- When you evaluate a deferrable function, any other deferrable function down the call tree will also be evaluated.
- You can include a call to a deferrable function in any Python expression and the result will be another deferred expression.

3.7.3 How a Deferred Expression Is Evaluated?

As discussed before, you can create a new deferred expression by calling a function whose definition is decorated by the `@sanity_function` or `@deferrable` decorator or by including an already deferred expression in any sort of arithmetic operation. When you call `evaluate` on a deferred expression, you trigger the evaluation of the whole subexpression tree. Here is how the evaluation process evolves:

A deferred expression object is merely a placeholder of the target function and its arguments at the moment you call it. Deferred expressions leverage also the Python's data model so as to capture all the binary and unary operators supported by the language. When you call `evaluate()` on a deferred expression object, the stored function will be called passing it the captured arguments. If any of the arguments is a deferred expression, it will be evaluated too. If the return value of the deferred expression is also a deferred expression, it will be evaluated as well.

This last property lets you call other deferrable functions from inside a deferrable function. Here is an example where we define two deferrable variations of the builtins `sum` and `len` and another deferrable function `avg()` that computes the average value of the elements of an iterable by calling our deferred builtin alternatives.

```
@sn.sanity_function
def dsum(iterable):
    return sum(iterable)

@sn.sanity_function
def dlen(iterable):
    return len(iterable)

@sn.sanity_function
def avg(iterable):
    return dsum(iterable) / dlen(iterable)
```

If you try to evaluate `avg()` with a list, you will get the expected result:

```
>>> avg([1, 2, 3, 4])
<reframe.core.deferrable._DeferredExpression object at 0x2b1288f54b70>
>>> evaluate(avg([1, 2, 3, 4]))
2.5
```

The return value of `evaluate(avg())` would normally be a deferred expression representing the division of the results of the other two deferrable functions. However, the evaluation mechanism detects that the return value is a deferred expression and it automatically triggers its evaluation, yielding the expected result. The following figure shows how the evaluation evolves for this particular example:

Fig. 2: Sequence diagram of the evaluation of the deferrable `avg()` function.

3.7.4 Implicit evaluation of a deferred expression

Although you can trigger the evaluation of a deferred expression at any time by calling `evaluate`, there are some cases where the evaluation is triggered implicitly:

- When you try to get the truthy value of a deferred expression by calling `bool` on it. This happens for example when you include a deferred expression in an `if` statement or as an argument to the `and`, `or`, `not` and `in` (`__contains__`) operators. The following example demonstrates this behavior:

```
>>> if avg([1, 2, 3, 4]) > 2:
...     print('hello')
...
hello
```

The expression `avg([1, 2, 3, 4]) > 2` is a deferred expression, but its evaluation is triggered from the Python interpreter by calling the `bool()` method on it, in order to evaluate the `if` statement. A similar example is the following that demonstrates the behaviour of the `in` operator:

```
>>> from reframe.utility.sanity import defer
>>> l = defer([1, 2, 3])
>>> l
<reframe.core.deferrable._DeferredExpression object at 0x2b1288f54cf8>
>>> evaluate(l)
[1, 2, 3]
>>> 4 in l
False
>>> 3 in l
True
```

The `defer` is simply a deferrable version of the identity function (a function that simply returns its argument). As expected, `l` is a deferred expression that evaluates to the `[1, 2, 3]` list. When we apply the `in` operator, the deferred expression is immediately evaluated.

Note: Python expands this expression into `bool(l.__contains__(3))`. Although `__contains__` is also defined as a deferrable function in `_DeferredExpression`, its evaluation is triggered by the `bool` builtin.

- When you try to iterate over a deferred expression by calling the `iter` function on it. This call happens implicitly by the Python interpreter when you try to iterate over a container. Here is an example:

```
>>> @sn.sanity_function
... def getlist(iterable):
...     ret = list(iterable)
...     ret += [1, 2, 3]
...     return ret
>>> getlist([1, 2, 3])
<reframe.core.deferrable._DeferredExpression object at 0x2b1288f54dd8>
>>> for x in getlist([1, 2, 3]):
...     print(x)
...
1
2
3
1
2
3
```

Simply calling `getlist()` will not execute anything and a deferred expression object will be returned. However, when you try to iterate over the result of this call, then the deferred expression will be evaluated immediately.

- When you try to call `str` on a deferred expression. This will be called by the Python interpreter every time you try to print this expression. Here is an example with the `getlist` deferrable function:

```
>>> print(getlist([1, 2, 3]))
[1, 2, 3, 1, 2, 3]
```

3.7.5 How to Write a Deferrable Function?

The answer is simple: like you would with any other normal function! We've done that already in all the examples we've shown in this documentation. A question that somehow naturally comes up here is whether you can call a deferrable function from within a deferrable function, since this doesn't make a lot of sense: after all, your function will be deferred anyway.

The answer is, yes. You can call other deferrable functions from within a deferrable function. Thanks to the implicit evaluation rules as well as the fact that the return value of a deferrable function is also evaluated if it is a deferred expression, you can write a deferrable function without caring much about whether the functions you call are themselves deferrable or not. However, you should be aware of passing mutable objects to deferrable functions. If these objects happen to change between the actual call and the implicit evaluation of the deferrable function, you might run into surprises. In any case, if you want the immediate evaluation of a deferrable function or expression, you can always do that by calling `evaluate` on it.

The following example demonstrates two different ways writing a deferrable function that checks the average of the elements of an iterable:

```
import reframe.utility.sanity as sn

@sn.sanity_function
def check_avg_with_deferrables(iterable):
    avg = sn.sum(iterable) / sn.len(iterable)
    return -1 if avg > 2 else 1

@sn.sanity_function
def check_avg_without_deferrables(iterable):
    avg = sum(iterable) / len(iterable)
    return -1 if avg > 2 else 1
```

```
>>> evaluate(check_avg_with_deferrables([1, 2, 3, 4]))
-1
>>> evaluate(check_avg_without_deferrables([1, 2, 3, 4]))
-1
```

The first version uses the `sum` and `len` functions from `reframe.utility.sanity`, which are deferrable versions of the corresponding builtins. The second version uses directly the builtin `sum` and `len` functions. As you can see, both of them behave in exactly the same way. In the version with the deferrables, `avg` is a deferred expression but it is evaluated by the `if` statement before returning.

Generally, inside a sanity function, it is a preferable to use the non-deferrable version of a function, if that exists, since you avoid the extra overhead and bookkeeping of the deferring mechanism.

3.7.6 Deferrable Sanity Functions

Normally, you will not have to implement your own sanity functions, since ReFrame provides already a variety of them. You can find the complete list of provided sanity functions [here](#).

Similarities and Differences with Generators

Python allows you to create functions that will be evaluated lazily. These are called `generator functions`. Their key characteristic is that instead of using the `return` keyword to return values, they use the `yield` keyword. I'm not going to go into the details of the generators, since there is plenty of documentation out there, so I will focus on the similarities and differences with our deferrable functions.

Similarities

- Both generators and our deferrables return an object representing the deferred expression when you call them.
- Both generators and deferrables may be evaluated explicitly or implicitly when they appear in certain expressions.
- When you try to iterate over a generator or a deferrable, you trigger its evaluation.

Differences

- You can include deferrables in any arithmetic expression and the result will be another deferrable expression. This is not true with generator functions, which will raise a `TypeError` in such cases or they will always evaluate to `False` if you include them in boolean expressions Here is an example demonstrating this:

```
>>> @sn.sanity_function
... def dsize(iterable):
...     print(len(iterable))
...     return len(iterable)
...
>>> def gsize(iterable):
...     print(len(iterable))
...     yield len(iterable)
...
>>> l = [1, 2]
>>> dsize(l)
<reframe.core.deferrable._DeferredExpression object at 0x2abc630abb38>
>>> gsize(l)
<generator object gsize at 0x2abc62a4bf10>
>>> expr = gsize(l) == 2
>>> expr
False
>>> expr = gsize(l) + 2
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
TypeError: unsupported operand type(s) for +: 'generator' and 'int'
>>> expr = dsize(l) == 2
>>> expr
<reframe.core.deferrable._DeferredExpression object at 0x2abc630abba8>
>>> expr = dsize(l) + 2
>>> expr
<reframe.core.deferrable._DeferredExpression object at 0x2abc630abc18>
```

Notice that you cannot include generators in expressions, whereas you can generate arbitrary expressions with deferrables.

- Generators are iterator objects, while deferred expressions are not. As a result, you can trigger the evaluation of a generator expression using the `next` builtin function. For a deferred expression you should use `evaluate` instead.

- A generator object is iterable, whereas a deferrable object will be iterable if and only if the result of its evaluation is iterable.

Note: Technically, a deferrable object is iterable, too, since it provides the `__iter__` method. That's why you can include it in iteration expressions. However, it delegates this call to the result of its evaluation.

Here is an example demonstrating this difference:

```
>>> for i in gsize(1): print(i)
...
2
2
>>> for i in dsize(1): print(i)
...
2
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "/users/karakasv/Devel/reframe/reframe/core/deferrable.py", line 73, in __
  ↪iter__
    return iter(self.evaluate())
TypeError: 'int' object is not iterable
```

Notice how the iteration works fine with the generator object, whereas with the deferrable function, the iteration call is delegated to the result of the evaluation, which is not an iterable, therefore yielding `TypeError`. Notice also, the printout of 2 in the iteration over the deferrable expression, which shows that it has been evaluated.

3.8 Running ReFrame

Before getting into any details, the simplest way to invoke ReFrame is the following:

```
./bin/reframe -c /path/to/checks -R --run
```

This will search recursively for test files in `/path/to/checks` and will start running them on the current system.

ReFrame's front-end goes through three phases:

1. Load tests
2. Filter tests
3. Act on tests

In the following, we will elaborate on these phases and the key command-line options controlling them. A detailed listing of all the command-line options grouped by phase is given by `./bin/reframe -h`.

3.8.1 Supported Actions

Even though an action is the last phase that the front-end goes through, we are listing it first since an action is always required. Currently there are only two available actions:

1. Listing of the selected checks
2. Execution of the selected checks

Listing of the regression tests

To retrieve a listing of the selected checks, you must specify the `-l` or `--list` options. This will provide a list with a brief description for each test containing only its name and the path to the file where it is defined. An example listing of checks is the following that lists all the tests found under the `tutorial/` folder:

```
./bin/reframe -c tutorial -l
```

The output looks like:

```
Command line: ./bin/reframe -c tutorial/ -l
Reframe version: 2.15-dev1
Launched by user: USER
Launched on host: daint103
Reframe paths
=====
  Check prefix      :
  Check search path : 'tutorial/'
  Stage dir prefix  : /path/to/reframe/stage/
  Output dir prefix : /path/to/reframe/output/
  Logging dir       : /path/to/reframe/logs
List of matched checks
=====
* Example5Test (found in /path/to/reframe/tutorial/example5.py)
* Example1Test (found in /path/to/reframe/tutorial/example1.py)
* Example4Test (found in /path/to/reframe/tutorial/example4.py)
* SerialTest   (found in /path/to/reframe/tutorial/example8.py)
* OpenMPTest   (found in /path/to/reframe/tutorial/example8.py)
* MPITest      (found in /path/to/reframe/tutorial/example8.py)
* OpenACCTest  (found in /path/to/reframe/tutorial/example8.py)
* CudaTest     (found in /path/to/reframe/tutorial/example8.py)
* Example3Test (found in /path/to/reframe/tutorial/example3.py)
* Example7Test (found in /path/to/reframe/tutorial/example7.py)
* Example6Test (found in /path/to/reframe/tutorial/example6.py)
* Example2aTest (found in /path/to/reframe/tutorial/example2.py)
* Example2bTest (found in /path/to/reframe/tutorial/example2.py)
Found 13 check(s).
```

You may also retrieve a listing with detailed information about the each check using the option `-L` or `--list-detailed`. The following example lists detailed information about the tutorial check:

```
Command line: ./bin/reframe -c tutorial/ -L
Reframe version: 2.18-dev2
Launched by user: USER
Launched on host: daint103
Reframe paths
=====
  Check prefix      :
  Check search path : 'tutorial/'
  Stage dir prefix  : /path/to/reframe/stage/
  Output dir prefix : /path/to/reframe/output/
  Logging dir       : /path/to/reframe/logs
List of matched checks
=====
* Example5Test (found in /path/to/reframe/tutorial/example5.py)
  - description: Matrix-vector multiplication example with CUDA
  - systems: daint:gpu
```

(continues on next page)

(continued from previous page)

```

- environments: PrgEnv-cray, PrgEnv-gnu, PrgEnv-pgi
- modules: cudatoolkit
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* Example1Test (found in /path/to/reframe/tutorial/example1.py)
- description: Simple matrix-vector multiplication example
- systems: *
- environments: *
- modules:
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* Example4Test (found in /path/to/reframe/tutorial/example4.py)
- description: Matrix-vector multiplication example with OpenACC
- systems: daint:gpu
- environments: PrgEnv-cray, PrgEnv-pgi
- modules: craype-accel-nvidia60
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* SerialTest (found in /path/to/reframe/tutorial/example8.py)
- description: Serial matrix-vector multiplication
- systems: *
- environments: *
- modules:
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* OpenMPTest (found in /path/to/reframe/tutorial/example8.py)
- description: OpenMP matrix-vector multiplication
- systems: *
- environments: PrgEnv-cray, PrgEnv-gnu, PrgEnv-intel, PrgEnv-pgi
- modules:
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* MPITest (found in /path/to/reframe/tutorial/example8.py)
- description: MPI matrix-vector multiplication
- systems: daint:gpu, daint:mc
- environments: PrgEnv-cray, PrgEnv-gnu, PrgEnv-intel, PrgEnv-pgi
- modules:
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* OpenACCTest (found in /path/to/reframe/tutorial/example8.py)
- description: OpenACC matrix-vector multiplication
- systems: daint:gpu
- environments: PrgEnv-cray, PrgEnv-pgi
- modules: craype-accel-nvidia60
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* CudaTest (found in /path/to/reframe/tutorial/example8.py)
- description: CUDA matrix-vector multiplication
- systems: daint:gpu
- environments: PrgEnv-gnu, PrgEnv-cray, PrgEnv-pgi

```

(continues on next page)

(continued from previous page)

```

- modules: cudatoolkit
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* Example3Test (found in /path/to/reframe/tutorial/example3.py)
- description: Matrix-vector multiplication example with MPI
- systems: daint:gpu, daint:mc
- environments: PrgEnv-cray, PrgEnv-gnu, PrgEnv-intel, PrgEnv-pgi
- modules:
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* Example7Test (found in /path/to/reframe/tutorial/example7.py)
- description: Matrix-vector multiplication (CUDA performance test)
- systems: daint:gpu
- environments: PrgEnv-gnu, PrgEnv-cray, PrgEnv-pgi
- modules: cudatoolkit
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* Example6Test (found in /path/to/reframe/tutorial/example6.py)
- description: Matrix-vector multiplication with L2 norm check
- systems: *
- environments: *
- modules:
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* Example2aTest (found in /path/to/reframe/tutorial/example2.py)
- description: Matrix-vector multiplication example with OpenMP
- systems: *
- environments: PrgEnv-cray, PrgEnv-gnu, PrgEnv-intel, PrgEnv-pgi
- modules:
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
* Example2bTest (found in /path/to/reframe/tutorial/example2.py)
- description: Matrix-vector multiplication example with OpenMP
- systems: *
- environments: PrgEnv-cray, PrgEnv-gnu, PrgEnv-intel, PrgEnv-pgi
- modules:
- task allocation: standard
- tags: tutorial
- maintainers: you-can-type-your-email-here
Found 13 check(s).

```

The detailed listing shows the description of the test, its supported systems and programming environments (* stands for any system or programming environment), the environment modules that it loads, its tags and its maintainers.

Warning: The list of modules showed in the detailed listing may not correspond to actual modules loaded by test, if the test customizes its behavior during the setup stage.

Note: New in version 2.15: Support for detailed listings. Standard listing using the `-l` option is now shorter.

Note: Changed in version 2.15: Test listing lists only tests supported by the current system. Previous versions were showing all the tests found.

Execution of the regression tests

To run the regression tests you should specify the *run* action though the `-r` or `--run` options.

Note: The listing action takes precedence over the execution, meaning that if you specify both `-l -r`, only the listing action will be performed.

```
./reframe.py -C tutorial/config/settings.py -c tutorial/example1.py -r
```

The output of the regression run looks like the following:

```
Command line: ./reframe.py -C tutorial/config/settings.py -c tutorial/example1.py -r
Reframe version: 2.13-dev0
Launched by user: USER
Launched on host: daint103
Reframe paths
=====
  Check prefix      :
  Check search path : 'tutorial/example1.py'
  Stage dir prefix  : /path/to/reframe/stage/
  Output dir prefix : /path/to/reframe/output/
  Logging dir       : /path/to/reframe/logs
[=====] Running 1 check(s)
[=====] Started on Sat May 26 00:34:34 2018

[-----] started processing Example1Test (Simple matrix-vector multiplication_
↪example)
[ RUN      ] Example1Test on daint:login using PrgEnv-cray
[ OK      ] Example1Test on daint:login using PrgEnv-cray
[ RUN      ] Example1Test on daint:login using PrgEnv-gnu
[ OK      ] Example1Test on daint:login using PrgEnv-gnu
[ RUN      ] Example1Test on daint:login using PrgEnv-intel
[ OK      ] Example1Test on daint:login using PrgEnv-intel
[ RUN      ] Example1Test on daint:login using PrgEnv-pgi
[ OK      ] Example1Test on daint:login using PrgEnv-pgi
[ RUN      ] Example1Test on daint:gpu using PrgEnv-cray
[ OK      ] Example1Test on daint:gpu using PrgEnv-cray
[ RUN      ] Example1Test on daint:gpu using PrgEnv-gnu
[ OK      ] Example1Test on daint:gpu using PrgEnv-gnu
[ RUN      ] Example1Test on daint:gpu using PrgEnv-intel
[ OK      ] Example1Test on daint:gpu using PrgEnv-intel
[ RUN      ] Example1Test on daint:gpu using PrgEnv-pgi
[ OK      ] Example1Test on daint:gpu using PrgEnv-pgi
[ RUN      ] Example1Test on daint:mc using PrgEnv-cray
[ OK      ] Example1Test on daint:mc using PrgEnv-cray
[ RUN      ] Example1Test on daint:mc using PrgEnv-gnu
[ OK      ] Example1Test on daint:mc using PrgEnv-gnu
[ RUN      ] Example1Test on daint:mc using PrgEnv-intel
[ OK      ] Example1Test on daint:mc using PrgEnv-intel
[ RUN      ] Example1Test on daint:mc using PrgEnv-pgi
```

(continues on next page)

(continued from previous page)

```
[          OK ] Example1Test on daint:mc using PrgEnv-pgi
[-----] finished processing Example1Test (Simple matrix-vector multiplication_
↪example)

[ PASSED ] Ran 12 test case(s) from 1 check(s) (0 failure(s))
[=====] Finished on Sat May 26 00:35:39 2018
```

3.8.2 Discovery of Regression Tests

When ReFrame is invoked, it tries to locate regression tests in a predefined path. By default, this path is the `<reframe-install-dir>/checks`. You can also retrieve this path as follows:

```
./bin/reframe -l | grep 'Check search path'
```

If the path line is prefixed with `(R)`, every directory in that path will be searched recursively for regression tests.

As described extensively in the “[ReFrame Tutorial](#)”, regression tests in ReFrame are essentially Python source files that provide a special function, which returns the actual regression test instances. A single source file may also provide multiple regression tests. ReFrame loads the python source files and tries to call this special function; if this function cannot be found, the source file will be ignored. At the end of this phase, the front-end will have instantiated all the tests found in the path.

You can override the default search path for tests by specifying the `-c` or `--checkpath` options. We have already done that already when listing all the tutorial tests:

```
./bin/reframe -c tutorial/ -l
```

ReFrame does not search recursively into directories specified with the `-c` option, unless you explicitly specify the `-R` or `--recurse` options.

The `-c` option completely overrides the default path. Currently, there is no option to prepend or append to the default regression path. However, you can build your own check path by specifying multiple times the `-c` option. The `-c` option accepts also regular files. This is very useful when you are implementing new regression tests, since it allows you to run only your test:

```
./bin/reframe -c /path/to/my/new/test.py -r
```

Important: The names of the loaded tests must be unique. Trying to load two or more tests with the same name will produce an error. You may ignore the error by using the `--ignore-check-conflicts` option. In this case, any conflicting test will not be loaded and a warning will be issued.

New in version 2.12.

3.8.3 Filtering of Regression Tests

At this phase you can select which regression tests should be run or listed. There are several ways to select regression tests, which we describe in more detail here:

Selecting tests by system

New in version 2.15.

By default, ReFrame always selects the tests that are supported by the current system. If you want to list the tests supported by a different system, you may achieve that by passing the `--system` option:

```
./bin/reframe --system=kesch -l
```

This example lists all the tests that are supported by the system named `kesch`. It is also possible to list only the tests that are supported by a specific system partition. The following example will list only the tests supported by the `login` partition of the `kesch` system:

```
./bin/reframe --system=kesch:login -l
```

Finally, in order to list all the tests found regardless of their supported systems, you should pass the `--skip-system-check` option:

```
./bin/reframe --skip-system-check -l
```

Selecting tests by programming environment

To select tests by the programming environment, use the `-p` or `--prgenv` options:

```
./bin/reframe -p PrgEnv-gnu -l
```

This will select all the checks that support the `PrgEnv-gnu` environment.

You can also specify multiple times the `-p` option, in which case a test will be selected if it support all the programming environments specified in the command line. For example the following will select all the checks that can run with both `PrgEnv-cray` and `PrgEnv-gnu` on the current system:

```
./bin/reframe -p PrgEnv-gnu -p PrgEnv-cray -l
```

If you are going to run a set of tests selected by programming environment, they will run only for the selected programming environment(s).

The `-p` option accepts also the Python [regular expression syntax](#). In fact, the argument to `-p` option is treated as a regular expression always. This means that the `-p PrgEnv-gnu` will match also tests that support a `PrgEnv-gnuXX` environment. If you would like to strictly select tests that support `PrgEnv-gnu` only and not `PrgEnv-gnuXX`, you should write `-p PrgEnv-gnu$`. As described above multiple `-p` options are AND-ed. Combining that with regular expressions can be quite powerful. For example, the following will select all tests that support programming environment `foo` and either `PrgEnv-gnu` or `PrgEnv-pgi`:

```
./bin/reframe -p foo -p 'PrgEnv-(gnu|pgi)' -l
```

Note: New in version 2.17.

The `-p` option recognizes regular expressions as arguments.

Selecting tests by tags

As we have seen in the “[ReFrame tutorial](#)”, every regression test may be associated with a set of tags. Using the `-t` or `--tag` option you can select the regression tests associated with a specific tag. For example the following will list all the tests that have a `maintenance` tag and can run on the current system:

```
./bin/reframe -t maintenance -l
```

Similarly to the `-p` option, you can chain multiple `-t` options together, in which case a regression test will be selected if it is associated with all the tags specified in the command line. The list of tags associated with a check can be viewed in the listing output when specifying the `-l` option.

Note: New in version 2.17.

The `-t` option recognizes regular expressions as arguments.

Selecting tests by name

It is possible to select or exclude tests by name through the `--name` or `-n` and `--exclude` or `-x` options. For example, you can select only the `Example7Test` from the tutorial as follows:

```
./bin/reframe -c tutorial/ -n Example7Test -l
```

```
Command line: ./bin/reframe -c tutorial/ -n Example7Test -l
Reframe version: 2.15-dev1
Launched by user: USER
Launched on host: daint103
Reframe paths
=====
  Check prefix      :
  Check search path : 'tutorial'
  Stage dir prefix  : /path/to/reframe/stage/
  Output dir prefix : /path/to/reframe/output/
  Logging dir       : /path/to/reframe/logs
List of matched checks
=====
 * Example7Test (found in /path/to/reframe/tutorial/example7.py)
Found 1 check(s).
```

Similarly, you can exclude this test by passing the `-x Example7Test` option:

```
Command line: ./bin/reframe -c tutorial -x Example7Test -l
Reframe version: 2.15-dev1
Launched by user: USER
Launched on host: daint103
Reframe paths
=====
  Check prefix      :
  Check search path : 'tutorial'
  Stage dir prefix  : /path/to/reframe/stage/
  Output dir prefix : /path/to/reframe/output/
  Logging dir       : /path/to/reframe/logs
List of matched checks
=====
 * Example5Test (found in /path/to/reframe/tutorial/example5.py)
 * Example1Test (found in /path/to/reframe/tutorial/example1.py)
 * Example4Test (found in /path/to/reframe/tutorial/example4.py)
 * SerialTest (found in /path/to/reframe/tutorial/example8.py)
 * OpenMPTest (found in /path/to/reframe/tutorial/example8.py)
 * MPITest (found in /path/to/reframe/tutorial/example8.py)
```

(continues on next page)

(continued from previous page)

```
* OpenACCTest (found in /path/to/reframe/tutorial/example8.py)
* CudaTest (found in /path/to/reframe/tutorial/example8.py)
* Example3Test (found in /path/to/reframe/tutorial/example3.py)
* Example6Test (found in /path/to/reframe/tutorial/example6.py)
* Example2aTest (found in /path/to/reframe/tutorial/example2.py)
* Example2bTest (found in /path/to/reframe/tutorial/example2.py)
Found 12 check(s).
```

Both `-n` and `-x` options can be chained, in which case either the tests that have any of the specified names are selected or excluded from running. They may also accept regular expressions as arguments.

Note: New in version 2.17: The `-n` and `-x` options recognize regular expressions as arguments. Chaining these options, e.g., `-n A -n B`, is equivalent to a regular expression that applies OR to the individual arguments, i.e., equivalent to `-n 'A|B'`.

3.8.4 Controlling the Execution of Regression Tests

There are several options for controlling the execution of regression tests. Keep in mind that these options will affect all the tests that will run with the current invocation. They are summarized below:

- `-A ACCOUNT, --account ACCOUNT`: Submit regression test jobs using `ACCOUNT`.
- `-P PART, --partition PART`: Submit regression test jobs in the *scheduler partition* `PART`.
- `--reservation RES`: Submit regression test jobs in reservation `RES`.
- `--nodelist NODELIST`: Run regression test jobs on the nodes specified in `NODELIST`.
- `--exclude-nodes NODELIST`: Do not run the regression test jobs on any of the nodes specified in `NODELIST`.
- `--job-option OPT`: Pass option `OPT` directly to the back-end job scheduler. This option *must* be used with care, since you may break the submission mechanism. All of the above job submission related options could be expressed with this option. For example, the `-n NODELIST` is equivalent to `--job-option='--nodelist=NODELIST'` for a Slurm job scheduler. If you pass an option that is already defined by the framework, the framework will *not* explicitly override it; this is up to scheduler. All extra options defined from the command line are appended to the automatically generated options in the generated batch script file. So if you redefine one of them, e.g., `--output` for the Slurm scheduler, it is up to the job scheduler on how to interpret multiple definitions of the same options. In this example, Slurm's policy is that later definitions of options override previous ones. So, in this case, you would override the standard output for all the submitted jobs!
- `--flex-alloc-tasks {all|idle|NUM}`: (Deprecated) Please use `--flex-alloc-nodes` instead.
- `--flex-alloc-nodes {all|idle|NUM}`: Automatically determine the number of nodes allocated for each test.
- `--force-local`: Force the local execution of the selected tests. No jobs will be submitted.
- `--skip-sanity-check`: Skip sanity checking phase.
- `--skip-performance-check`: Skip performance verification phase.
- `--strict`: Force strict performance checking. Some tests may set their `strict_check` attribute to `False` (see “[Reference Guide](#)”) in order to just let their performance recorded but not yield an error. This option overrides this behavior and forces all tests to be strict.

- `--skip-system-check`: Skips the system check and run the selected tests even if they do not support the current system. This option is sometimes useful when you need to quickly verify if a regression test supports a new system.
- `--skip-prgenv-check`: Skips programming environment check and run the selected tests for even if they do not support a programming environment. This option is useful when you need to quickly verify if a regression check supports another programming environment. For example, if you know that a tests supports only `PrgEnv-cray` and you need to check if it also works with `PrgEnv-gnu`, you can test is as follows:

```
./bin/reframe -c /path/to/my/check.py -p PrgEnv-gnu --skip-prgenv-check -r
```

- `--max-retries NUM`: Specify the maximum number of times a failed regression test may be retried (default: 0).

3.8.5 Generating a Performance Report

If you are running performance tests, you may instruct ReFrame to produce a performance report at the end using the `-performance-report` command-line options. The performance report is printed after the output of the regression tests and has the following format:

```
PERFORMANCE REPORT
-----
Check1
- system:partition
  - PrgEnv1
    * num_tasks: <num_tasks>
    * perf_variable1: <value> <units>
    * perf_variable2: <value> <units>
    * ...
  - PrgEnv2
    * num_tasks: <num_tasks>
    * perf_variable1: <value> <units>
    * perf_variable2: <value> <units>
    * ...
-----
Check2
- system:partition
  - PrgEnv1
    * num_tasks: <num_tasks>
    * perf_variable1: <value> <units>
    * perf_variable2: <value> <units>
    * ...
  - PrgEnv2
    * num_tasks: <num_tasks>
    * perf_variable1: <value> <units>
    * perf_variable2: <value> <units>
    * ...
-----
```

The number of tasks and the achieved performance values are listed by system partition and programming environment for each performance test that has run. Performance variables are the variables collected through the `reframe.core.pipeline.RegressionTest.perf_patterns` attribute.

The following command will run the CUDA matrix-vector multiplication example from the [tutorial](#) and will produce a performance report:

```
./bin/reframe -C tutorial/config/settings.py -c tutorial/example7.py -r --performance-
↳report
```

```
Command line: ./bin/reframe -C tutorial/config/settings.py -c tutorial/example7.py -r
↳--performance-report
Reframe version: 2.20-dev2
Launched by user: USER
Launched on host: daint101
Reframe paths
=====
Check prefix      :
Check search path : 'example7.py'
Stage dir prefix  : /path/to/reframe/stage/
Output dir prefix : /path/to/reframe/output/
Perf. logging prefix : /path/to/reframe/perflogs
[=====] Running 1 check(s)
[=====] Started on Thu Oct 24 17:46:55 2019

[-----] started processing Example7Test (Matrix-vector multiplication (CUDA
↳performance test))
[ RUN      ] Example7Test on daint:gpu using PrgEnv-cray
[ OK       ] Example7Test on daint:gpu using PrgEnv-cray
[ RUN      ] Example7Test on daint:gpu using PrgEnv-gnu
[ OK       ] Example7Test on daint:gpu using PrgEnv-gnu
[ RUN      ] Example7Test on daint:gpu using PrgEnv-pgi
[ OK       ] Example7Test on daint:gpu using PrgEnv-pgi
[-----] finished processing Example7Test (Matrix-vector multiplication (CUDA
↳performance test))

[ PASSED  ] Ran 3 test case(s) from 1 check(s) (0 failure(s))
[=====] Finished on Thu Oct 24 17:47:34 2019
=====
PERFORMANCE REPORT
-----
Example7Test
- daint:gpu
  - PrgEnv-cray
    * num_tasks: 1
    * perf: 49.403965 Gflop/s
  - PrgEnv-gnu
    * num_tasks: 1
    * perf: 50.093877 Gflop/s
  - PrgEnv-pgi
    * num_tasks: 1
    * perf: 50.549009 Gflop/s
-----
```

For completeness, we show here the corresponding section from the `Example7Test`, so that the connection between the test's code and the output becomes clear:

```
self.perf_patterns = {
    'perf': sn.extractsingle(r'Performance:\s+(?P<Gflops>\S+) Gflop/s',
                           self.stdout, 'Gflops', float)
}
self.reference = {
    'daint:gpu': {
        'perf': (50.0, -0.1, 0.1, 'Gflop/s'),
```

(continues on next page)

(continued from previous page)

```
}
}
```

If you are writing a benchmark, it is often the case that you will run it in an unknown system, where you don't have any reference value. Normally, if ReFrame cannot find a reference for the system it is running on, it will complain and mark the test as a failure. However, you may right your test in such a way, that it allows it to run successfully on any new system. To achieve this, simply insert a “catch-all” * entry in the `reframe.core.pipeline.RegressionTest.reference` attribute:

```
self.reference = {
    '*': {
        'perf_var1': (0, None, None, 'units'),
        'perf_var2': (0, None, None, 'units')
        ...
    }
}
```

The performance test will always pass on new systems and you may use the `--performance-report` option for getting the actual performance values.

Note: The performance report should not be confused with *performance logging*. It is simply a way of quickly visualizing the performance results and is useful for interactive testing. Performance logging, if configured, occurs independently of the performance report and is meant for keeping performance data over time. Its formatting facilitates parsing and it should be used for later analysis of the performance data obtained.

3.8.6 Configuring ReFrame Directories

ReFrame uses two basic directories during the execution of tests:

1. The stage directory

- Each regression test is executed in a “sandbox”; all of its resources (source files, input data etc.) are copied over to a stage directory (if the directory preexists, it will be wiped out) and executed from there. This will also be the working directory for the test.

2. The output directory

- After a regression test finishes some important files will be copied from the stage directory to the output directory (if the directory preexists, it will be wiped out). By default these are the standard output, standard error and the generated job script file. A regression test may also specify to keep additional files.

By default, these directories are placed under a common prefix, which defaults to `..`. The rest of the directories are organized as follows:

- Stage directory: `${prefix}/stage/<timestamp>`
- Output directory: `${prefix}/output/<timestamp>`

You can optionally append a timestamp directory component to the above paths (except the logs directory), by using the `--timestamp` option. This options takes an optional argument to specify the timestamp format. The default time format is `%FT%T`, which results into timestamps of the form `2017-10-24T21:10:29`.

You can override either the default global prefix or any of the default individual directories using the corresponding options.

- `--prefix DIR`: set prefix to DIR.

- `--output DIR`: set output directory to DIR.
- `--stage DIR`: set stage directory to DIR.

The stage and output directories are created only when you run a regression test. However you can view the directories that will be created even when you do a listing of the available checks with the `-l` option. This is useful if you want to check the directories that ReFrame will create.

```
./bin/reframe -C tutorial/config/settings.py --prefix /foo -l
```

```
Command line: ./bin/reframe -C tutorial/config/settings.py --prefix /foo -l
Reframe version: 2.13-dev0
Launched by user: USER
Launched on host: daint103
Reframe paths
=====
    Check prefix      : /path/to/reframe
(R) Check search path : 'checks/'
    Stage dir prefix  : /foo/stage/
    Output dir prefix : /foo/output/
    Perf. logging prefix : /Users/karakasv/Repositories/reframe/logs
List of matched checks
=====
Found 0 check(s).
```

You can also define different default directories per system by specifying them in the [site configuration settings file](#). The command line options, though, take always precedence over any default directory.

3.8.7 Logging

From version 2.4 onward, ReFrame supports logging of its actions. ReFrame creates two files inside the current working directory every time it is run:

- `reframe.out`: This file stores the output of a run as it was printed in the standard output.
- `reframe.log`: This file stores more detailed of information on ReFrame's actions.

By default, the output in `reframe.log` looks like the following:

```
2018-05-26T00:30:39] info: reframe: [ RUN          ] Example7Test on daint:gpu using_
↳PrgEnv-cray
[2018-05-26T00:30:39] debug: Example7Test: entering stage: setup
[2018-05-26T00:30:39] debug: Example7Test: loading environment for the current_
↳partition
[2018-05-26T00:30:39] debug: Example7Test: executing OS command: modulecmd python_
↳show daint-gpu
[2018-05-26T00:30:39] debug: Example7Test: executing OS command: modulecmd python_
↳load daint-gpu
[2018-05-26T00:30:39] debug: Example7Test: loading test's environment
[2018-05-26T00:30:39] debug: Example7Test: executing OS command: modulecmd python_
↳show PrgEnv-cray
[2018-05-26T00:30:39] debug: Example7Test: executing OS command: modulecmd python_
↳unload PrgEnv-gnu
[2018-05-26T00:30:39] debug: Example7Test: executing OS command: modulecmd python_
↳load PrgEnv-cray
[2018-05-26T00:30:39] debug: Example7Test: executing OS command: modulecmd python_
↳show cudatoolkit
[2018-05-26T00:30:39] debug: Example7Test: executing OS command: modulecmd python_
↳load cudatoolkit
```

(continues on next page)

(continued from previous page)

```

[2018-05-26T00:30:39] debug: Example7Test: setting up paths
[2018-05-26T00:30:40] debug: Example7Test: setting up the job descriptor
[2018-05-26T00:30:40] debug: Example7Test: job scheduler backend: local
[2018-05-26T00:30:40] debug: Example7Test: setting up performance logging
[2018-05-26T00:30:40] debug: Example7Test: entering stage: compile
[2018-05-26T00:30:40] debug: Example7Test: copying /path/to/reframe/tutorial/src to
↳stage directory (/path/to/reframe/stage/gpu/Example7Test/PrgEnv-cray)
[2018-05-26T00:30:40] debug: Example7Test: symlinking files: []
[2018-05-26T00:30:40] debug: Example7Test: Staged sourcepath: /path/to/reframe/stage/
↳gpu/Example7Test/PrgEnv-cray/example_matrix_vector_multiplication_cuda.cu
[2018-05-26T00:30:40] debug: Example7Test: executing OS command: nvcc -O3 -I/path/to/
↳reframe/stage/gpu/Example7Test/PrgEnv-cray /path/to/reframe/stage/gpu/Example7Test/
↳PrgEnv-cray/e
xample_matrix_vector_multiplication_cuda.cu -o /path/to/reframe/stage/gpu/
↳Example7Test/PrgEnv-cray/./Example7Test
[2018-05-26T00:30:40] debug: Example7Test: compilation stdout:

[2018-05-26T00:30:40] debug: Example7Test: compilation stderr:
nvcc warning : The 'compute_20', 'sm_20', and 'sm_21' architectures are deprecated,
↳and may be removed in a future release (Use -Wno-deprecated-gpu-targets to suppress
↳warning).

[2018-05-26T00:30:40] debug: Example7Test: compilation finished
[2018-05-26T00:30:40] debug: Example7Test: entering stage: run
[2018-05-26T00:30:40] debug: Example7Test: executing OS command: sbatch /path/to/
↳reframe/stage/gpu/Example7Test/PrgEnv-cray/Example7Test_daint_gpu_PrgEnv-cray.sh
[2018-05-26T00:30:40] debug: Example7Test: spawned job (jobid=746641)
[2018-05-26T00:30:40] debug: Example7Test: entering stage: wait
[2018-05-26T00:30:40] debug: Example7Test: executing OS command: sacct -S 2018-05-26 -
↳P -j 746641 -o jobid,state,exitcode
[2018-05-26T00:30:40] debug: Example7Test: job state not matched (stdout follows)
JobID|State|ExitCode

[2018-05-26T00:30:41] debug: Example7Test: executing OS command: sacct -S 2018-05-26 -
↳P -j 746641 -o jobid,state,exitcode
[2018-05-26T00:30:44] debug: Example7Test: executing OS command: sacct -S 2018-05-26 -
↳P -j 746641 -o jobid,state,exitcode
[2018-05-26T00:30:47] debug: Example7Test: executing OS command: sacct -S 2018-05-26 -
↳P -j 746641 -o jobid,state,exitcode
[2018-05-26T00:30:47] debug: Example7Test: spawned job finished
[2018-05-26T00:30:47] debug: Example7Test: entering stage: sanity
[2018-05-26T00:30:47] debug: Example7Test: entering stage: performance
[2018-05-26T00:30:47] debug: Example7Test: entering stage: cleanup
[2018-05-26T00:30:47] debug: Example7Test: copying interesting files to output
↳directory
[2018-05-26T00:30:47] debug: Example7Test: removing stage directory
[2018-05-26T00:30:47] info: reframe: [      OK ] Example7Test on daint:gpu using
↳PrgEnv-cray

```

Each line starts with a timestamp, the level of the message (info, debug etc.), the context in which the framework is currently executing (either `reframe` or the name of the current test and, finally, the actual message.

Every time ReFrame is run, both `reframe.out` and `reframe.log` files will be rewritten. However, you can ask ReFrame to copy them to the output directory before exiting by passing it the `--save-log-files` option.

Configuring Logging

You can configure several aspects of logging in ReFrame and even how the output will look like. ReFrame's logging mechanism is built upon Python's [logging](#) framework adding extra logging levels and more formatting capabilities.

Logging in ReFrame is configured by the `logging_config` variable in the `reframe/settings.py` file. The default configuration looks as follows:

```
logging_config = {
    'level': 'DEBUG',
    'handlers': [
        {
            'type': 'file',
            'name': 'reframe.log',
            'level': 'DEBUG',
            'format': '[%(asctime)s] %(levelname)s: '
                    '%(check_info)s: %(message)s',
            'append': False,
        },

        # Output handling
        {
            'type': 'stream',
            'name': 'stdout',
            'level': 'INFO',
            'format': '%(message)s'
        },
        {
            'type': 'file',
            'name': 'reframe.out',
            'level': 'INFO',
            'format': '%(message)s',
            'append': False,
        }
    ]
}
```

Note that this configuration dictionary is not the same as the one used by Python's logging framework. It is a simplified version adapted to the needs of ReFrame.

The `logging_config` dictionary has two main key entries:

- `level` (default: `'INFO'`): This is the lowest level of messages that will be passed down to the different log record handlers. Any message with a lower level than that, it will be filtered out immediately and will not be passed to any handler. ReFrame defines the following logging levels with a decreasing severity: `CRITICAL`, `ERROR`, `WARNING`, `INFO`, `VERBOSE` and `DEBUG`. Note that the level name is *not* case sensitive in ReFrame.
- `handlers`: A list of log record handlers that are attached to ReFrame's logging mechanism. You can attach as many handlers as you like. For example, by default ReFrame uses three handlers: (a) a handler that logs debug information into `reframe.log`, (b) a handler that controls the actual output of the framework to the standard output, which does not print any debug messages, and (c) a handler that writes the same output to a file `reframe.out`.

Each handler is configured by another dictionary that holds its properties as string key/value pairs. For standard ReFrame logging there are currently two types of handlers, which recognize different properties.

Note: New syntax for handlers is introduced. The old syntax is still valid, but users are advised to update their logging configuration to the new syntax.

Changed in version 2.13.

Common Log Handler Attributes

All handlers accept the following set of attributes (keys) in their configuration:

- `type`: (required) the type of the handler. There are several types of handlers used for logging in ReFrame. Some of them are only relevant for performance logging:
 1. `file`: a handler that writes log records in file.
 2. `stream`: a handler that writes log records in a file stream.
 3. `syslog`: a handler that sends log records to Unix syslog.
 4. `filelog`: a handler for writing performance logs (relevant only for *performance logging*).
 5. `graylog`: a handler for sending performance logs to a Graylog server (relevant only for *performance logging*).
- `level`: (default: `DEBUG`) The lowest level of log records that this handler can process.
- `format` (default: `'%(message)s'`): Format string for the printout of the log record. ReFrame supports all the [format strings](#) from Python's logging library and provides the following additional ones:
 - `check_environ`: The programming environment a test is currently executing for.
 - `check_info`: Print live information of the currently executing check. By default this field has the form `<check_name> on <current_partition> using <current_environment>`. It can be configured on a per test basis by overriding the `info` method of a specific regression test.
 - `check_jobid`: Prints the job or process id of the job or process associated with the currently executing regression test. If a job or process is not yet created, `-1` will be printed.
 - `check_job_completion_time`: *[new in 2.21]* The completion time of the job spawned by this regression test. This timestamp will be formatted according to `datefmt` (see below). The accuracy of the timestamp depends on the backend scheduler. The `slurm` scheduler backend relies on job accounting and returns the actual termination time of the job. The rest of the backends report as completion time the moment when the framework realizes that the spawned job has finished. In this case, the accuracy depends on the execution policy used. If tests are executed with the serial execution policy, this is close to the real completion time, but if the asynchronous execution policy is used, it can differ significantly. If the job completion time cannot be retrieved, `None` will be printed.
 - `check_name`: Prints the name of the regression test on behalf of which ReFrame is currently executing. If ReFrame is not in the context of regression test, `reframe` will be printed.
 - `check_num_tasks`: The number of tasks assigned to the regression test.
 - `check_outputdir`: The output directory associated with the currently executing test.
 - `check_partition`: The system partition where this test is currently executing.
 - `check_stagedir`: The stage directory associated with the currently executing test.
 - `check_system`: The host system where this test is currently executing.
 - `check_tags`: The tags associated with this test.
 - `osuser`: The name of the OS user running ReFrame.
 - `osgroup`: The group name of the OS user running ReFrame.
 - `version`: The ReFrame version.

- `datefmt` (default: `'%FT%T'`) The format that will be used for outputting timestamps (i.e., the `%(asctime)s` field). Acceptable formats must conform to standard library's `time.strftime()` function.

Caution: The `testcase_name` logging attribute is replaced with the `check_info`, which is now also configurable

Changed in version 2.10.

File log handlers

In addition to the common log handler attributes, file log handlers accept the following:

- `name`: (required) The name of the file where log records will be written.
- `append` (default: `False`) Controls whether ReFrame should append to this file or not.
- `timestamp` (default: `None`): Append a timestamp to this log filename. This property may accept any date format that is accepted also by the `datefmt` property. If the name of the file is `filename.log` and this attribute is set to `True`, the resulting log file name will be `filename_<timestamp>.log`.

Stream log handlers

In addition to the common log handler attributes, file log handlers accept the following:

- `name`: (default `stdout`) The symbolic name of the log stream to use. Available values: `stdout` for standard output and `stderr` for standard error.

Syslog log handler

In addition to the common log handler attributes, file log handlers accept the following:

- `socktype`: The type of socket where the handler will send log records to. There are two socket types:
 1. `udp`: (default) This opens a UDP datagram socket.
 2. `tcp`: This opens a TCP stream socket.
- `facility`: (default: `user`) The Syslog facility to send records to. The list of supported facilities can be found [here](#).
- `address`: (required) The address where the handler will connect to. This can either be of the form `<host>:<port>` or simply a path that refers to a Unix domain socket.

Note: New in version 2.17.

Performance Logging

ReFrame supports an additional logging facility for recording performance values, in order to be able to keep historical performance data. This is configured by the `perf_logging_config` variables, whose syntax is the same as for the `logging_config`:

```

perf_logging_config = {
    'level': 'DEBUG',
    'handlers': [
        {
            'type': 'filelog',
            'prefix': '%(check_system)s/%(check_partition)s',
            'level': 'INFO',
            'format': (
                '%(asctime)s|reframe %(version)s|'
                '%(check_info)s|jobid=%(check_jobid)s|'
                '%(check_perf_var)s=%(check_perf_value)s|'
                'ref=%(check_perf_ref)s|'
                '(l=%(check_perf_lower_thres)s, '
                'u=%(check_perf_upper_thres)s)|'
                '%(check_perf_unit)s'
            ),
            'append': True
        }
    ]
}

```

Performance logging introduces two new log record handlers, specifically designed for this purpose.

File-based Performance Logging

The type of this handler is `filelog` and logs the performance of a regression test in one or more files. The attributes of this handler are the following:

- `prefix`: This is the directory prefix (usually dynamic) where the performance logs of a test will be stored. This attribute accepts any of the check-specific formatting placeholders described *above*. This allows you to create dynamic paths based on the current system, partition and/or programming environment a test executes. This dynamic prefix is appended to the “global” performance log directory prefix, configurable through the `--perflogdir` option or the `perflogdir` attribute of the [system configuration](#). The default configuration of ReFrame for performance logging (shown in the previous listing) generates the following files:

```

{PERFLOG_PREFIX}/
  system1/
    partition1/
      test_name.log
    partition2/
      test_name.log
    ...
  system2/
    ...

```

A log file, named after the test’s name, is generated in different directories, which are themselves named after the system and partition names that this test has run on. The `PERFLOG_PREFIX` will have the value of `--perflogdir` option, if specified, otherwise it will default to `{REFRAME_PREFIX}/perflogs`. You can always check its value by looking into the paths printed by ReFrame at the beginning of its output:

```

Command line: ./reframe.py --prefix=/foo --system=generic -l
Reframe version: 2.13-dev0
Launched by user: USER
Launched on host: HOSTNAME
Reframe paths

```

(continues on next page)

(continued from previous page)

```

=====
  Check prefix      : /Users/karakasv/Repositories/reframe
(R) Check search path : 'checks/'
  Stage dir prefix  : /foo/stage/
  Output dir prefix  : /foo/output/
  Perf. logging prefix : /foo/perflogs
List of matched checks
=====
Found 0 check(s).

```

- **format:** The syntax of this attribute is the same as of the standard logging facility, except that it adds a couple more performance-specific formatting placeholders:
 - `check_perf_lower_thres`: The lower threshold of the difference from the reference value expressed as a fraction of the reference.
 - `check_perf_upper_thres`: The upper threshold of the difference from the reference value expressed as a fraction of the reference.
 - `check_perf_ref`: The reference performance value of a certain performance variable.
 - `check_perf_value`: The performance value obtained by this test for a certain performance variable.
 - `check_perf_var`: The name of the [performance variable](#), whose value is logged.
 - `check_perf_unit`: The unit of measurement for the measured performance variable, if specified in the corresponding tuple of the `reframe.core.pipeline.RegressionTest.reference` attribute.

Note: Changed in version 2.20: Support for logging `num_tasks` in performance logs was added.

Using the default performance log format, the resulting log entries look like the following:

```

2019-10-23T13:46:05|reframe 2.20-dev2|Example7Test on daint:gpu using PrgEnv-
↪cray|jobid=813559|num_tasks=1|perf=49.681565|ref=50.0 (l=-0.1, u=0.1)|Gflop/s
2019-10-23T13:46:27|reframe 2.20-dev2|Example7Test on daint:gpu using PrgEnv-
↪gnu|jobid=813560|num_tasks=1|perf=50.737651|ref=50.0 (l=-0.1, u=0.1)|Gflop/s
2019-10-23T13:46:48|reframe 2.20-dev2|Example7Test on daint:gpu using PrgEnv-
↪pgi|jobid=813561|num_tasks=1|perf=50.720164|ref=50.0 (l=-0.1, u=0.1)|Gflop/s

```

The interpretation of the performance values depends on the individual tests. The above output is from the CUDA performance test we presented in the [tutorial](#), so the value refers to the achieved Gflop/s.

Performance Logging Using Graylog

The type of this handler is `graylog` and it logs performance data to a [Graylog](#) server. Graylog is a distributed enterprise log management service. An example configuration of such a handler is the following:

```

{
  'type': 'graylog',
  'host': 'my.graylog.server',
  'port': 12345,
  'level': 'INFO',
  'format': (
    '%(asctime)s|reframe %(version)s|'
    '%(check_info)s|jobid=%(check_jobid)s|'
  )
}

```

(continues on next page)

(continued from previous page)

```

    'num_tasks=%(check_num_tasks)s|'
    '%(check_perf_var)s=%(check_perf_value)s|'
    'ref=%(check_perf_ref)s|'
    '(l=%(check_perf_lower_thres)s,|'
    'u=%(check_perf_upper_thres)s|'
),
'extras': {
    'facility': 'reframe',
}
},

```

This handler introduces three new attributes:

- `host`: (required) The Graylog server that accepts the log messages.
- `port`: (required) The port where the Graylog server accepts connections.
- `extras`: (optional) A set of optional user attributes to be passed with each log record to the server. These may depend on the server configuration.

This log handler uses internally `pygelf`, so this Python module must be available, otherwise this log handler will be ignored. `GELF` is a format specification for log messages that are sent over the network. The ReFrame's `graylog` handler sends log messages in JSON format using an HTTP POST request to the specified host and port. More details on this log format may be found [here](#).

Adjusting verbosity of output

ReFrame's output is handled by a logging mechanism. In fact, as revealed in the corresponding configuration entry (see *Configuring Logging*), a specific logging handler takes care of printing ReFrame's message in the standard output. One way to change the verbosity level of the output is by explicitly setting the value of the `level` key in the configuration of the output handler. Alternatively, you may increase the verbosity level from the command line by chaining the `-v` or `--verbose` option. Every time `-v` is specified, the next verbosity level will be selected for the output. For example, if the initial level of the output handler is set to `INFO` (in the configuration file), specifying `-v` twice will make ReFrame spit out all `DEBUG` messages.

New in version 2.16: `-v` and `--verbose` options are added.

3.8.8 Asynchronous Execution of Regression Checks

From version 2.4, ReFrame supports asynchronous execution of regression tests. This execution policy can be enabled by passing the option `--exec-policy=async` to the command line. The default execution policy is `serial` which enforces a sequential execution of the selected regression tests. The asynchronous execution policy parallelizes only the `running phase` of the tests. The rest of the phases remain sequential.

A limit of concurrent jobs (pending and running) may be `configured` for each virtual system partition. As soon as the concurrency limit of a partition is reached, ReFrame will hold the execution of new regression tests until a slot is released in that partition.

When executing in asynchronous mode, ReFrame's output differs from the sequential execution. The final result of the tests will be printed at the end and additional messages may be printed to indicate that a test is held. Here is an example output of ReFrame using asynchronous execution policy:

```

Command line: ./bin/reframe -C tutorial/config/settings.py -c tutorial/ --exec-
→policy=async -r
Reframe version: 2.13-dev0

```

(continues on next page)

(continued from previous page)

```

Launched by user: USER
Launched on host: daint103
Reframe paths
=====
    Check prefix      :
    Check search path : 'tutorial/'
    Stage dir prefix  : /path/to/reframe/stage/
    Output dir prefix : /path/to/reframe/output/
    Logging dir       : /path/to/reframe/logs
[=====] Running 13 check(s)
[=====] Started on Sat May 26 00:48:03 2018

[-----] started processing Example1Test (Simple matrix-vector multiplication_
↳example)
[ RUN      ] Example1Test on daint:login using PrgEnv-cray
[ RUN      ] Example1Test on daint:login using PrgEnv-gnu
[ RUN      ] Example1Test on daint:login using PrgEnv-intel
[ RUN      ] Example1Test on daint:login using PrgEnv-pgi
[ RUN      ] Example1Test on daint:gpu using PrgEnv-cray
[ RUN      ] Example1Test on daint:gpu using PrgEnv-gnu
[ RUN      ] Example1Test on daint:gpu using PrgEnv-intel
[ RUN      ] Example1Test on daint:gpu using PrgEnv-pgi
[ RUN      ] Example1Test on daint:mc using PrgEnv-cray
[ RUN      ] Example1Test on daint:mc using PrgEnv-gnu
[ RUN      ] Example1Test on daint:mc using PrgEnv-intel
[ RUN      ] Example1Test on daint:mc using PrgEnv-pgi
[-----] finished processing Example1Test (Simple matrix-vector multiplication_
↳example)

[-----] started processing Example2aTest (Matrix-vector multiplication example_
↳with OpenMP)
[ RUN      ] Example2aTest on daint:login using PrgEnv-cray
[ RUN      ] Example2aTest on daint:login using PrgEnv-gnu
[ RUN      ] Example2aTest on daint:login using PrgEnv-intel
[ RUN      ] Example2aTest on daint:login using PrgEnv-pgi
[ RUN      ] Example2aTest on daint:gpu using PrgEnv-cray
[ RUN      ] Example2aTest on daint:gpu using PrgEnv-gnu
[ RUN      ] Example2aTest on daint:gpu using PrgEnv-intel
[ RUN      ] Example2aTest on daint:gpu using PrgEnv-pgi
[ RUN      ] Example2aTest on daint:mc using PrgEnv-cray
[ RUN      ] Example2aTest on daint:mc using PrgEnv-gnu
[ RUN      ] Example2aTest on daint:mc using PrgEnv-intel
[ RUN      ] Example2aTest on daint:mc using PrgEnv-pgi
[-----] finished processing Example2aTest (Matrix-vector multiplication example_
↳with OpenMP)
<output omitted>
[-----] waiting for spawned checks to finish
[ OK ] MPITest on daint:gpu using PrgEnv-pgi
[ OK ] MPITest on daint:gpu using PrgEnv-gnu
[ OK ] OpenMPTest on daint:mc using PrgEnv-pgi
[ OK ] OpenMPTest on daint:mc using PrgEnv-gnu
[ OK ] OpenMPTest on daint:gpu using PrgEnv-pgi
[ OK ] OpenMPTest on daint:gpu using PrgEnv-gnu
<output omitted>
[ OK ] Example1Test on daint:login using PrgEnv-cray
[ OK ] MPITest on daint:mc using PrgEnv-cray
[ OK ] MPITest on daint:gpu using PrgEnv-cray

```

(continues on next page)

(continued from previous page)

```

[      OK ] OpenMPTest on daint:mc using PrgEnv-cray
[      OK ] OpenMPTest on daint:gpu using PrgEnv-cray
[      OK ] SerialTest on daint:login using PrgEnv-pgi
[      OK ] MPITest on daint:mc using PrgEnv-gnu
[      OK ] OpenMPTest on daint:mc using PrgEnv-intel
[      OK ] OpenMPTest on daint:login using PrgEnv-gnu
[      OK ] OpenMPTest on daint:gpu using PrgEnv-intel
[      OK ] MPITest on daint:gpu using PrgEnv-intel
[      OK ] CudaTest on daint:gpu using PrgEnv-gnu
[      OK ] OpenACCTest on daint:gpu using PrgEnv-pgi
[      OK ] MPITest on daint:mc using PrgEnv-intel
[      OK ] CudaTest on daint:gpu using PrgEnv-cray
[      OK ] MPITest on daint:mc using PrgEnv-pgi
[      OK ] OpenACCTest on daint:gpu using PrgEnv-cray
[      OK ] CudaTest on daint:gpu using PrgEnv-pgi
[-----] all spawned checks have finished

[ PASSED ] Ran 101 test case(s) from 13 check(s) (0 failure(s))
[=====] Finished on Sat May 26 00:52:02 2018

```

The asynchronous execution policy may provide significant overall performance benefits for run-only regression tests. For compile-only and normal tests that require a compilation, the execution time will be bound by the total compilation time of the test.

3.8.9 Manipulating modules

New in version 2.11.

Note: Changed in version 2.19: Module self loops are now allowed in *module mappings*.

ReFrame allows you to change the modules loaded by a regression test on-the-fly without having to edit the regression test file. This feature is extremely useful when you need to quickly test a newer version of a module, but it also allows you to completely decouple the module names used in your regression tests from the real module names in a system, thus making your test even more portable. This is achieved by defining *module mappings*.

There are two ways to pass module mappings to ReFrame. The first is to use the `--map-module` command-line option, which accepts a module mapping. For example, the following line maps the module `test_module` to the module `real_module`:

```
--map-module='test_module: real_module'
```

In this case, whenever ReFrame is asked to load `test_module`, it will load `real_module`. Any string without spaces may be accepted in place of `test_module` and `real_module`. You can also define multiple module mappings at once by repeating the `--map-module`. If more than one mapping is specified for the same module, then the last mapping will take precedence. It is also possible to map a single module to more than one target. This can be done by listing the target modules separated by spaces in the order that they should be loaded. In the following example, ReFrame will load `real_module0` and `real_module1` whenever the `test_module` is encountered:

```
--map-module 'test_module: real_module0 real_module1'
```

The second way of defining mappings is by listing them on a file, which you can then pass to ReFrame through the command-line option `--module-mappings`. Each line on the file corresponds to the definition of a mapping for a single module. The syntax of the individual mappings in the file is the same as with the option `--map-module` and

the same rules apply regarding repeated definitions. Text starting with # is considered a comment and is ignored until the end of line is encountered. Empty lines are ignored. The following block shows an example of module mapping file:

```
module-1: module-1a # an inline comment
module-2: module-2a module-2b module-2c

# This is a full line comment
module-4: module-4a module-4b
```

If both `--map-module` and `--module-mappings` are passed, ReFrame will first create a mapping from the definitions on the file and it will then process the definitions passed with the `--map-module` options. As usual, later definitions will override the former.

A final note on module mappings. Module mappings can be arbitrarily deep as long as they do not form a cycle. In this case, ReFrame will issue an error (denoting the offending cyclic dependency). For example, suppose having the following mapping file:

```
cuda toolkit: foo
foo: bar
bar: foobar
foobar: cuda toolkit
```

If you now try to run a test that loads the module *cuda toolkit*, the following error will be yielded:

```
-----
FAILURE INFO for Example7Test
 * System partition: daint:gpu
 * Environment: PrgEnv-gnu
 * Stage directory: None
 * Job type: batch job (id=-1)
 * Maintainers: ['you-can-type-your-email-here']
 * Failing phase: setup
 * Reason: caught framework exception: module cyclic dependency: cuda toolkit->foo->
↳bar->foobar->cuda toolkit
-----
```

On the other hand, module mappings containing self loops are allowed. In the following example, ReFrame will load both `module-1` and `module-2` whenever the `module-1` is encountered:

```
--map-module 'module-1: module-1 module-2'
```

3.8.10 Controlling the Flexible Node Allocation

New in version 2.15.

Note: Changed in version 2.21: Flexible task allocation is now based on number of nodes.

Warning: The command line option `--flex-alloc-tasks` is now deprecated, you should use `--flex-alloc-nodes` instead.

ReFrame can automatically set the number of tasks of a particular test, if its `num_tasks` attribute is set to a value ≤ 0 . By default, ReFrame will spawn such a test on all the idle nodes of the current system partition. This behavior

can be adjusted using the `--flex-alloc-nodes` command line option. This option accepts three values:

1. `idle`: (default) In this case, ReFrame will set the number of tasks to the number of idle nodes of the current logical partition multiplied by the `num_tasks_per_node` attribute of the particular test.
2. `all`: In this case, ReFrame will set the number of tasks to the number of all the nodes of the current logical partition multiplied by the `num_tasks_per_node` attribute of the particular test.
3. Any positive integer: In this case, ReFrame will set the number of tasks to the given value multiplied by the `num_tasks_per_node` attribute of the particular test.

The flexible allocation of number of nodes takes into account any additional logical constraint imposed by the command line options affecting the job allocation, such as `--partition`, `--reservation`, `--nodelist`, `--exclude-nodes` and `--job-option` (if the scheduler option passed to the latter imposes a restriction). Notice that ReFrame will issue an error if the resulting number of nodes is zero.

For example, using the following options would run a flexible test on all the nodes of reservation `foo` except the nodes `n0[1-5]`:

```
--flex-alloc-nodes=all --reservation=foo --exclude-nodes=n0[1-5]
```

Note: Flexible node allocation is supported only for the Slurm scheduler backend.

Warning: Test cases resulting from flexible ReFrame tests may not be run using the asynchronous execution policy, because the nodes satisfying the required criteria will be allocated for the first test case, causing all subsequent ones to fail.

3.8.11 Testing non-default Cray Programming Environments

New in version 2.20.

Cray machines provide a set of compilers, scientific software libraries and an MPI implementation that is optimized for the Cray hardware. These comprise the Cray Programming Environment (PE). All the functionality of the PE is structured around the default versions (or modules) of the libraries. If a non-default library or a non-default PE are to be tested, users have to do the following after having loaded all of their Cray modules:

```
export LD_LIBRARY_PATH=$CRAY_LD_LIBRARY_PATH:$LD_LIBRARY_PATH
```

In order to test a non-default Cray Programming Environment with ReFrame you have to pass the `--non-default-craype`. This will cause ReFrame to export `LD_LIBRARY_PATH` as shown above. Here is an example that shows how to test a non-default Cray PE with ReFrame:

```
module load cdt/19.08
reframe <options> --non-default-craype -r
```

3.9 Use Cases

3.9.1 ReFrame Usage at CSCS

The ReFrame framework has been in production at [CSCS](#) since December 2016. We use it to test not only [Piz Daint](#), but almost all our systems that we provide to users.

We have two large sets of regression tests:

- production tests and
- maintenance tests.

Tags are used to mark these categories and a regression test may belong to both of them. Production tests are run daily to monitor the sanity of the system and its performance. All performance tests log their performance values. The performance over time of certain applications are monitored graphically using [Grafana](#).

The total set of our regression tests comprises 172 individual tests, from which 153 are marked as production tests. Some of them are eligible to run on both the multicore and hybrid partitions of the system, whereas others are meant to run only on the login nodes. Depending on the test, multiple programming environments might be tried. In total, 448 test cases are run from the 153 regression tests on all the system partitions. The following Table summarizes the production regression tests.

The set of maintenance regression tests is much more limited to decrease the downtime of the system. The regression suite runs at the beginning of the maintenance session and just before returning the machine to the users, so that we can ensure that the user experience is at least at the level before the system was taken down. The maintenance set of tests comprises application performance tests, some GPU library performance checks, Slurm checks and some POSIX filesystem checks.

The porting of the regression suite to the [MeteoSwiss](#) production system [Piz Kesch](#), using ReFrame was almost trivial. The new system entry was added in the framework's configuration file describing the different partitions together with a new redefined `PrgEnv-gnu` environment to use different compiler wrappers. Porting the regression tests of interest was also a straightforward process. In most of the cases, adding just the corresponding system partitions to the `valid_systems` variables and adjusting accordingly the `valid_prog_environs` was enough.

ReFrame really focuses on abstracting away all the gory details from the regression test description, hence letting the user to concentrate solely on the logic of his test. A bit of this effect can be seen in the following Table where the total amount of lines of code (loc) of the regression tests written in the previous shell script-based solution and ReFrame is shown. We also present a snapshot of the first public release of ReFrame ([v2.2](#)).

Maintenance Burden	Shell-Script Based	ReFrame (May 2017)	ReFrame (Nov 2017)
Total tests	179	122	172
Total size of tests	14635 loc	2985 loc	4493 loc
Avg. test file size	179 loc	93 loc	87 loc
Avg. effective test size	179 loc	25 loc	25 loc

The difference in the total amount of regression test code is dramatic. From the 15K lines of code of the old shell script based regression testing suite, ReFrame tests use only 3K lines of code (first release) achieving a higher coverage.

Note: The higher test count of the older suite refers to test cases, i.e., running the same test for different programming environments, whereas for ReFrame the counts do not account for this.

Each regression test file in ReFrame is 80–90 loc on average. However, each regression test file may contain or generate more than one related tests, thus leading to the effective decrease of the line count per test to only 25 loc.

Separating the logical description of a regression test from all the unnecessary implementation details contributes significantly in the ease of writing and maintaining new regression tests with ReFrame.

3.10 About ReFrame

3.10.1 What Is ReFrame?

ReFrame is a framework developed by CSCS to facilitate the writing of regression tests that check the sanity of HPC systems. Its main goal is to allow users to write their own regression tests without having to deal with all the details of setting up the environment for the test, querying the status of their job, managing the output of the job and looking for sanity and/or performance results. Users should be concerned only about the logical requirements of their tests. This allows users' regression checks to be maintained and adapted to new systems easily.

The user describes his test in a simple Python class and the framework takes care of all the details of the low-level interaction with the system. The framework is structured in such a way that with a basic knowledge of Python and minimal coding a user can write a regression test, which will be able to run out-of-the-box on a variety of systems and programming environments.

Writing regression tests in a high-level language, such as Python, allows users to take advantage of the language's higher expressiveness and bigger capabilities compared to classical shell scripting, which is the norm in HPC testing. This could lead to a more manageable code base of regression tests with significantly reduced maintenance costs.

3.10.2 ReFrame's Goals

When designing the framework we have set three major goals:

Productivity The writer of a regression test should focus only on the logical structure and requirements of the test and should not need to deal with any of the low level details of interacting with the system, e.g., how the environment of the test is loaded, how the associated job is created and has its status checked, how the output parsing is performed etc.

Portability Configuring the framework to support new systems and system configurations should be easy and should not affect the existing tests. Also, adding support of a new system in a regression test should require minimal adjustments.

Robustness and ease of use The new framework must be stable enough and easy to use by non-advanced users. When the system needs to be returned to users outside normal working hours the personnel in charge should be able to run the regression suite and verify the sanity of the system with a minimal involvement.

3.10.3 Why ReFrame?

HPC systems are highly complex systems in all levels of integration; from the physical infrastructure up to the software stack provided to the users. A small change in any of these levels could have an impact on the stability or the performance of the system perceived by the end users. It is of crucial importance, therefore, not only to make sure that the system is in a sane condition after every maintenance before handing it off to users, but also to monitor its performance during production, so that possible problems are detected early enough and the quality of service is not compromised.

Regression testing can provide a reliable way to ensure the stability and the performance requirements of the system, provided that sufficient tests exist that cover a wide aspect of the system's operations from both the operators' and users' point of view. However, given the complexity of HPC systems, writing and maintaining regression tests can be a very time consuming task. A small change in system configuration or deployment may require adapting hundreds of regression tests at the same time. Similarly, porting a test to a different system may require significant effort if the new system's configuration is substantially different than that of the system that it was originally written for.

ReFrame was designed to help HPC support teams to easily write tests that

- monitor the impact of changes to the system that would affect negatively the users,

- monitor system performance,
- monitor system stability and
- guarantee quality of service.

And also decrease the amount of time and resources required to

- write and maintain regression tests and
- port regression tests to other HPC systems.

3.11 Reference Guide

This page provides a reference guide of the ReFrame API for writing regression tests covering all the relevant details. Internal data structures and APIs are covered only to the extent that might be helpful to the final user of the framework.

3.11.1 Regression test classes and related utilities

`reframe.core.decorators.parameterized_test` (*inst)

Class decorator for registering multiple instantiations of a test class.

The decorated class must derive from `reframe.core.pipeline.RegressionTest`. This decorator is also available directly under the `reframe` module.

Parameters `inst` – The different instantiations of the test. Each instantiation argument may be either a sequence or a mapping.

New in version 2.13.

Note: This decorator does not instantiate any test. It only registers them. The actual instantiation happens during the loading phase of the test.

`reframe.core.decorators.simple_test` (cls)

Class decorator for registering parameterless tests with ReFrame.

The decorated class must derive from `reframe.core.pipeline.RegressionTest`. This decorator is also available directly under the `reframe` module.

New in version 2.13.

`reframe.core.decorators.required_version` (*versions)

Class decorator for specifying the required ReFrame versions for the following test.

If the test is not compatible with the current ReFrame version it will be skipped.

Parameters `versions` – A list of ReFrame version specifications that this test is allowed to run. A version specification string can have one of the following formats:

1. `VERSION`: Specifies a single version.
2. `{OP}VERSION`, where `{OP}` can be any of `>`, `>=`, `<`, `<=`, `==` and `!=`. For example, the version specification string `'>=2.15'` will only allow the following test to be loaded only by ReFrame 2.15 and higher. The `==VERSION` specification is the equivalent of `VERSION`.
3. `V1 . . V2`: Specifies a range of versions.

You can specify multiple versions with this decorator, such as `@required_version('2.13', '>=2.16')`, in which case the test will be selected if *any* of the versions is satisfied, even if the versions specifications are conflicting.

New in version 2.13.

`reframe.core.decorators.require_deps` (*func*)

Denote that the decorated test method will use the test dependencies.

The arguments of the decorated function must be named after the dependencies that the function intends to use. The decorator will bind the arguments to a partial realization of the `reframe.core.pipeline.RegressionTest.getdep()` function, such that conceptually the new function arguments will be the following:

```
new_arg = functools.partial(getdep, orig_arg_name)
```

The converted arguments are essentially functions accepting a single argument, which is the target test's programming environment.

This decorator is also directly available under the `reframe` module.

New in version 2.21.

`reframe.core.decorators.run_before` (*stage*)

Run the decorated function before the specified pipeline stage.

The decorated function must be a method of a regression test.

New in version 2.20.

`reframe.core.decorators.run_after` (*stage*)

Run the decorated function after the specified pipeline stage.

The decorated function must be a method of a regression test.

New in version 2.20.

class `reframe.core.pipeline.RegressionTest`

Bases: `object`

Base class for regression tests.

All regression tests must eventually inherit from this class. This class provides the implementation of the pipeline phases that the regression test goes through during its lifetime.

Parameters

- **name** – The name of the test. If `None`, the framework will try to assign a unique and human-readable name to the test.
- **prefix** – The directory prefix of the test. If `None`, the framework will set it to the directory containing the test file.

Note:

The `name` and `prefix` arguments are just maintained for backward compatibility to the old (prior to 2.13) syntax of regression tests. Users are advised to use the new simplified syntax for writing regression tests. Refer to the *ReFrame Tutorial* for more information.

This class is also directly available under the top-level `reframe` module.

Changed in version 2.13.

build_system

The build system to be used for this test. If not specified, the framework will try to figure it out automatically based on the value of `sourcepath`.

This field may be set using either a string referring to a concrete build system class name (see [build systems](#)) or an instance of `reframe.core.buildsystems.BuildSystem`. The former is the recommended way.

Type `str` or `reframe.core.buildsystems.BuildSystem`.

Default `None`.

New in version 2.14.

check_performance()

The performance checking phase of the regression test pipeline.

Raises `reframe.core.exceptions.SanityError` – If the performance check fails.

check_sanity()

The sanity checking phase of the regression test pipeline.

Raises `reframe.core.exceptions.SanityError` – If the sanity check fails.

cleanup(remove_files=False)

The cleanup phase of the regression test pipeline.

Parameters `remove_files` – If `True`, the stage directory associated with this test will be removed.

compile()

The compilation phase of the regression test pipeline.

Raises `reframe.core.exceptions.ReframeError` – In case of errors.

compile_wait()

Wait for compilation phase to finish.

New in version 2.13.

container_platform

The container platform to be used for launching this test.

If this field is set, the test will run inside a container using the specified container runtime. Container-specific options must be defined additionally after this field is set:

```
self.container_platform = 'Singularity'
self.container_platform.image = 'docker://ubuntu:18.04'
self.container_platform.commands = ['cat /etc/os-release']
```

If this field is set, `executable` and `executable_opts` attributes are ignored. The container platform's `commands` will be used instead. For more information on the container platform support, see the [tutorial](#) and the [reference guide](#).

Type `str` or `reframe.core.containers.ContainerPlatform`.

Default `None`.

New in version 2.20.

current_environ

The programming environment that the regression test is currently executing with.

This is set by the framework during the `setup()` phase.

Type `reframe.core.environments.Environment`.

current_partition

The system partition the regression test is currently executing on.

This is set by the framework during the `setup()` phase.

Type `reframe.core.systems.SystemPartition`.

current_system

The system the regression test is currently executing on.

This is set by the framework during the initialization phase.

Type `reframe.core.runtime.HostSystem`.

depends_on (*target*, *how*=2, *subdeps*=None)

Add a dependency to *target* in this test.

Parameters

- **target** – The name of the target test.
- **how** – How the dependency should be mapped in the test cases space. This argument can accept any of the three constants `DEPEND_EXACT`, `DEPEND_BY_ENV` (default), `DEPEND_FULLLY`.
- **subdeps** – An adjacency list representation of how this test's test cases depend on those of the target test. This is only relevant if `how == DEPEND_EXACT`. The value of this argument is a dictionary having as keys the names of this test's supported programming environments. The values are lists of the programming environments names of the target test that this test's test cases will depend on. In the following example, this test's E0 programming environment case will depend on both E0 and E1 test cases of the target test T0, but its E1 case will depend only on the E1 test case of T0:

```
self.depends_on('T0', how=rfm.DEPEND_EXACT,
                subdeps={'E0': ['E0', 'E1'], 'E1': ['E1']})
```

For more details on how test dependencies work in ReFrame, please refer to [How Test Dependencies Work In ReFrame](#).

New in version 2.21.

descr

A detailed description of the test.

Type `str`

Default `self.name`

exclusive_access

Specify whether this test needs exclusive access to nodes.

Type `boolean`

Default `False`

executable

The name of the executable to be launched during the run phase.

Type `str`

Default `os.path.join('.', self.name)`

executable_opts

List of options to be passed to the `executable`.

Type List[str]

Default []

extra_resources

Extra resources for this test.

This field is for specifying custom resources needed by this test. These resources are defined in the *configuration* of a system partition. For example, assume that two additional resources, named `gpu` and `datawarp`, are defined in the configuration file as follows:

```
'resources': {
  'gpu': [
    '--gres=gpu:{num_gpus_per_node}'
  ],
  'datawarp': [
    '#DW jobdw capacity={capacity}',
    '#DW stage_in source={stagein_src}'
  ]
}
```

A regression test then may instantiate the above resources by setting the `extra_resources` attribute as follows:

```
self.extra_resources = {
  'gpu': {'num_gpus_per_node': 2}
  'datawarp': {
    'capacity': '100GB',
    'stagein_src': '/foo'
  }
}
```

The generated batch script (for Slurm) will then contain the following lines:

```
#SBATCH --gres=gpu:2
#DW jobdw capacity=100GB
#DW stage_in source=/foo
```

Notice that if the resource specified in the configuration uses an alternative directive prefix (in this case `#DW`), this will replace the standard prefix of the backend scheduler (in this case `#SBATCH`)

If the resource name specified in this variable does not match a resource name in the partition configuration, it will be simply ignored. The `num_gpus_per_node` attribute translates internally to the `_rfm_gpu` resource, so that setting `self.num_gpus_per_node = 2` is equivalent to the following:

```
self.extra_resources = {'_rfm_gpu': {'num_gpus_per_node': 2}}
```

Type Dict[str, Dict[str, object]]

Default {}

Note: New in version 2.8.

Changed in version 2.9.

A new more powerful syntax was introduced that allows also custom job script directive prefixes.

getdep (*target*, *environ=None*)

Retrieve the test case of a target dependency.

This is a low-level method. The `@require_deps` decorators should be preferred.

Parameters

- **target** – The name of the target dependency to be retrieved.
- **environ** – The name of the programming environment that will be used to retrieve the test case of the target test. If `None`, `RegressionTest.current_environ` will be used.

New in version 2.21.

info ()

Provide live information of a running test.

This method is used by the front-end to print the status message during the test's execution. This function is also called to provide the message for the `check_info` logging attribute. By default, it returns a message reporting the test name, the current partition and the current programming environment that the test is currently executing on.

Returns a string with an informational message about this test

Note: When overriding this method, you should pay extra attention on how you use the `RegressionTest`'s attributes, because this method may be called at any point of the test's lifetime.

New in version 2.10.

is_local ()

Check if the test will execute locally.

A test executes locally if the `local` attribute is set or if the current partition's scheduler does not support job submission.

job

The job descriptor associated with this test.

This is set by the framework during the `setup()` phase.

Type `reframe.core.schedulers.Job`.

keep_files

List of files to be kept after the test finishes.

By default, the framework saves the standard output, the standard error and the generated shell script that was used to run this test.

These files will be copied over to the framework's output directory during the `cleanup()` phase.

Directories are also accepted in this field.

Relative path names are resolved against the stage directory.

Type `List[str]`

Default `[]`

local

Always execute this test locally.

Type `boolean`

Default False

logger

A logger associated with this test.

You can use this logger to log information for your test.

maintainers

List of people responsible for this test.

When the test fails, this contact list will be printed out.

Type List[str]

Default []

modules

List of modules to be loaded before running this test.

These modules will be loaded during the `setup()` phase.

Type List[str]

Default []

name

The name of the test.

Type string that can contain any character except /

num_cpus_per_task

Number of CPUs per task required by this test.

Ignored if None.

Type integral or None

Default None

num_gpus_per_node

Number of GPUs per node required by this test.

Type integral

Default 0

num_tasks

Number of tasks required by this test.

If the number of tasks is set to a number ≤ 0 , ReFrame will try to flexibly allocate the number of tasks, based on the command line option `--flex-alloc-nodes`. A negative number is used to indicate the minimum number of tasks required for the test. In this case the minimum number of tasks is the absolute value of the number, while Setting `num_tasks` to 0 is equivalent to setting it to `-num_tasks_per_node`.

Type integral

Default 1

Note: Changed in version 2.15: Added support for flexible allocation of the number of tasks according to the `--flex-alloc-tasks` command line option (see [Flexible node allocation](#)) if the number of tasks is set to 0.

Changed in version 2.16: Negative `num_tasks` is allowed for specifying the minimum number of required tasks by the test.

Changed in version 2.21: Flexible node allocation is now controlled by the `--flex-alloc-nodes` command line option (see [Flexible node allocation](#))

num_tasks_per_core

Number of tasks per core required by this test.

Ignored if `None`.

Type `integral` or `None`

Default `None`

num_tasks_per_node

Number of tasks per node required by this test.

Ignored if `None`.

Type `integral` or `None`

Default `None`

num_tasks_per_socket

Number of tasks per socket required by this test.

Ignored if `None`.

Type `integral` or `None`

Default `None`

outputdir

The output directory of the test.

This is set during the `setup()` phase.

New in version 2.13.

Type `str`.

perf_patterns

Patterns for verifying the performance of this test.

Refer to the *ReFrame Tutorial* for concrete usage examples.

If set to `None`, no performance checking will be performed.

Type A dictionary with keys of type `str` and deferrable expressions (i.e., the result of a *sanity function*) as values. `None` is also allowed.

Default `None`

poll()

Poll the test's state.

Returns

`True` if the associated job has finished, `False` otherwise.

If no job descriptor is yet associated with this test, `True` is returned.

Raises `reframe.core.exceptions.ReframeError` – In case of errors.

post_run

List of shell commands to execute after launching this job.

See *pre_run* for a more detailed description of the semantics.

Type `List[str]`

Default `[]`

Note: New in version 2.10.

postbuild_cmd

List of shell commands to be executed after a successful compilation.

These commands are executed during the compilation phase and from inside the stage directory. **Each entry in the list spawns a new shell.**

Type `List[str]`

Default `[]`

pre_run

List of shell commands to execute before launching this job.

These commands do not execute in the context of ReFrame. Instead, they are emitted in the generated job script just before the actual job launch command.

Type `List[str]`

Default `[]`

Note: New in version 2.10.

prebuild_cmd

List of shell commands to be executed before compiling.

These commands are executed during the compilation phase and from inside the stage directory. **Each entry in the list spawns a new shell.**

Type `List[str]`

Default `[]`

prefix

The prefix directory of the test.

Type `str`.

readonly_files

List of files or directories (relative to the `sourcesdir`) that will be symlinked in the stage directory and not copied.

You can use this variable to avoid copying very large files to the stage directory.

Type `List[str]`

Default `[]`

reference

The set of reference values for this test.

The reference values are specified as a scoped dictionary keyed on the performance variables defined in `perf_patterns` and scoped under the system/partition combinations. The reference itself is a three- or four-tuple that contains the reference value, the lower and upper thresholds and, optionally, the measurement unit. An example follows:

```
self.reference = {
    'sys0:part0': {
        'perfvar0': (50, -0.1, 0.1, 'Gflop/s'),
        'perfvar1': (20, -0.1, 0.1, 'GB/s')
    },
    'sys0:part1': {
        'perfvar0': (100, -0.1, 0.1, 'Gflop/s'),
        'perfvar1': (40, -0.1, 0.1, 'GB/s')
    }
}
```

Type A scoped dictionary with system names as scopes or None

Default {}

run()

The run phase of the regression test pipeline.

This call is non-blocking. It simply submits the job associated with this test and returns.

sanity_patterns

Refer to the *ReFrame Tutorial* for concrete usage examples.

If set to None, a sanity error will be raised during sanity checking.

Type A deferrable expression (i.e., the result of a *sanity function*) or None

Default None

Note: Changed in version 2.9: The default behaviour has changed and it is now considered a sanity failure if this attribute is set to None.

If a test doesn't care about its output, this must be stated explicitly as follows:

```
self.sanity_patterns = sn.assert_found(r'.*', self.stdout)
```

setup (*partition*, *environ*, ****job_opts**)

The setup phase of the regression test pipeline.

Parameters

- **partition** – The system partition to set up this test for.
- **environ** – The environment to set up this test for.
- **job_opts** – Options to be passed through to the backend scheduler. When overriding this method users should always pass through `job_opts` to the base class method.

Raises `reframe.core.exceptions.ReframeError` – In case of errors.

sourcepath

The path to the source file or source directory of the test.

It must be a path relative to the `sourcesdir`, pointing to a subfolder or a file contained in `sourcesdir`. This applies also in the case where `sourcesdir` is a Git repository.

If it refers to a regular file, this file will be compiled using the *SingleSource* build system. If it refers to a directory, ReFrame will try to infer the build system to use for the project and will fall back in using the *Make* build system, if it cannot find a more specific one.

Type `str`

Default `''`

sourcesdir

The directory containing the test's resources.

This directory may be specified with an absolute path or with a path relative to the location of the test. Its contents will always be copied to the stage directory of the test.

This attribute may also accept a URL, in which case ReFrame will treat it as a Git repository and will try to clone its contents in the stage directory of the test.

If set to `None`, the test has no resources and no action is taken.

Type `str` or `None`

Default `'src'`

Note: Changed in version 2.9: Allow `None` values to be set also in regression tests with a compilation phase

Changed in version 2.10: Support for Git repositories was added.

stagedir

The stage directory of the test.

This is set during the `setup()` phase.

Type `str`.

stderr

The name of the file containing the standard error of the test.

This is set during the `setup()` phase.

This attribute is evaluated lazily, so it can be used inside sanity expressions.

Type `str`.

stdout

The name of the file containing the standard output of the test.

This is set during the `setup()` phase.

This attribute is evaluated lazily, so it can be used inside sanity expressions.

Type `str`.

strict_check

Mark this test as a strict performance test.

If a test is marked as non-strict, the performance checking phase will always succeed, unless the `--strict` command-line option is passed when invoking ReFrame.

Type `boolean`

Default `True`

tags

Set of tags associated with this test.

This test can be selected from the frontend using any of these tags.

Type `Set[str]`

Default an empty set

time_limit

Time limit for this test.

Time limit is specified as a three-tuple in the form (hh, mm, ss), with hh >= 0, 0 <= mm <= 59 and 0 <= ss <= 59. If set to None, no time limit will be set. The default time limit of the system partition's scheduler will be used.

Type tuple[int]

Default (0, 10, 0)

Note: Changed in version 2.15.

This attribute may be set to None.

use_multithreading

Specify whether this tests needs simultaneous multithreading enabled.

Ignored if None.

Type boolean or None

Default None

valid_prog_environs

List of programming environments supported by this test.

If * is in the list then all programming environments are supported by this test.

Type List[str]

Default []

Note: Changed in version 2.12: Programming environments can now be specified using wildcards.

Changed in version 2.17: Support for wildcards is dropped.

valid_systems

List of systems supported by this test. The general syntax for systems is <sysname>[:<partname>].

Type List[str]

Default []

variables

Environment variables to be set before running this test.

These variables will be set during the *setup()* phase.

Type Dict[str, str]

Default {}

wait()

Wait for this test to finish.

Raises `reframe.core.exceptions.ReframeError` – In case of errors.

class `reframe.core.pipeline.RunOnlyRegressionTest`

Bases: `reframe.core.pipeline.RegressionTest`

Base class for run-only regression tests.

This class is also directly available under the top-level `reframe` module.

compile()

The compilation phase of the regression test pipeline.

This is a no-op for this type of test.

compile_wait()

Wait for compilation phase to finish.

This is a no-op for this type of test.

run()

The run phase of the regression test pipeline.

The resources of the test are copied to the stage directory and the rest of execution is delegated to the `RegressionTest.run()`.

class `reframe.core.pipeline.CompileOnlyRegressionTest`

Bases: `reframe.core.pipeline.RegressionTest`

Base class for compile-only regression tests.

These tests are by default local and will skip the run phase of the regression test pipeline.

The standard output and standard error of the test will be set to those of the compilation stage.

This class is also directly available under the top-level `reframe` module.

run()

The run stage of the regression test pipeline.

Implemented as no-op.

setup(*partition, environ, **job_opts*)

The setup stage of the regression test pipeline.

Similar to the `RegressionTest.setup()`, except that no job descriptor is set up for this test.

stderr

The name of the file containing the standard error of the test.

This is set during the `setup()` phase.

This attribute is evaluated lazily, so it can be used inside sanity expressions.

Type `str`.

stdout

The name of the file containing the standard output of the test.

This is set during the `setup()` phase.

This attribute is evaluated lazily, so it can be used inside sanity expressions.

Type `str`.

wait()

Wait for this test to finish.

Implemented as no-op

`reframe.core.pipeline.DEPEND_EXACT = 1`

Constant to be passed as the `how` argument of the `RegressionTest.depends_on()` method. It denotes that test case dependencies will be explicitly specified by the user.

This constant is directly available under the `reframe` module.

```
reframe.core.pipeline.DEPEND_BY_ENV = 2
```

Constant to be passed as the `how` argument of the `RegressionTest.depends_on()` method. It denotes that the test cases of the current test will depend only on the corresponding test cases of the target test that use the same programming environment.

This constant is directly available under the `reframe` module.

```
reframe.core.pipeline.DEPEND_FULLY = 3
```

Constant to be passed as the `how` argument of the `RegressionTest.depends_on()` method. It denotes that each test case of this test depends on all the test cases of the target test.

This constant is directly available under the `reframe` module.

3.11.2 Environments and Systems

```
class reframe.core.environments.Environment (name, modules=[], variables=[])
```

Bases: `object`

This class abstracts away an environment to run regression tests.

It is simply a collection of modules to be loaded and environment variables to be set when this environment is loaded by the framework.

details()

Return a detailed description of this environment.

is_loaded

True if this environment is loaded, False otherwise.

modules

The modules associated with this environment.

Type `list of str`

name

The name of this environment.

Type `str`

variables

The environment variables associated with this environment.

Type dictionary of `str` keys/values.

```
class reframe.core.environments.ProgEnvironment (name, modules=[], variables={},
                                                cc='cc', cxx='CC', fn='fn',
                                                nvcc='nvcc', cppflags=None,
                                                cflags=None, cxxflags=None,
                                                fflags=None, ldflags=None,
                                                **kwargs)
```

Bases: `reframe.core.environments.Environment`

A class representing a programming environment.

This type of environment adds also attributes for setting the compiler and compilation flags.

If compilation flags are set to `None` (the default, if not set otherwise in ReFrame's `configuration`), they are not passed to the `make` invocation.

If you want to disable completely the propagation of the compilation flags to the `make` invocation, even if they are set, you should set the `propagate` attribute to `False`.

cc
The C compiler of this programming environment.

Type `str`

cflags
The C compiler flags of this programming environment.

Type `str` or `None`

cppflags
The preprocessor flags of this programming environment.

Type `str` or `None`

cxx
The C++ compiler of this programming environment.

Type `str` or `None`

cxxflags
The C++ compiler flags of this programming environment.

Type `str` or `None`

details ()
Return a detailed description of this environment.

fflags
The Fortran compiler flags of this programming environment.

Type `str` or `None`

ftn
The Fortran compiler of this programming environment.

Type `str` or `None`

ldflags
The linker flags of this programming environment.

Type `str` or `None`

`reframe.core.environments.load (*environs)`
Load environments in the current Python context.

Returns a tuple containing a snapshot of the environment at entry to this function and a list of shell commands required to load `environs`.

`reframe.core.environments.snapshot ()`
Create an environment snapshot

class `reframe.core.environments.temp_environment (modules=[], variables=[])`
Bases: `object`

Context manager to temporarily change the environment.

class `reframe.core.systems.System (name, descr=None, hostnames=[], partitions=[], preload_env=None, prefix='.', stagedir=None, output_dir=None, perflogdir=None, resourcesdir='.', modules_system=None)`

Bases: `object`

A representation of a system inside ReFrame.

descr

The description of this system.

hostnames

The hostname patterns associated with this system.

modules_system

The modules system name associated with this system.

name

The name of this system.

outputdir

The ReFrame output directory prefix associated with this system.

partitions

All the system partitions associated with this system.

perflogdir

The ReFrame log directory prefix associated with this system.

prefix

The ReFrame prefix associated with this system.

preload_environ

The environment to load whenever ReFrame runs on this system.

Note: New in version 2.19.

resourcesdir

Global resources directory for this system.

You may use this directory for storing large resource files of your regression tests. See [here](#) on how to configure this.

Type `str`

stagedir

The ReFrame stage directory prefix associated with this system.

class `reframe.core.systems.SystemPartition` (*name*, *descr=None*, *scheduler=None*,
launcher=None, *access=[]*, *environs=[]*,
resources={}, *local_env=None*, *max_jobs=1*)

Bases: `object`

A representation of a system partition inside ReFrame.

This class is immutable.

descr

A detailed description of this partition.

fullname

Return the fully-qualified name of this partition.

The fully-qualified name is of the form `<parent-system-name>:<partition-name>`.

Type `str`

launcher

The type of the backend launcher of this partition.

Returns a subclass of `reframe.core.launchers.JobLauncher`.

Note: New in version 2.8.

name

The name of this partition.

Type *str*

scheduler

The type of the backend scheduler of this partition.

Returns a subclass of `reframe.core.schedulers.Job`.

Note: Changed in version 2.8.

Prior versions returned a string representing the scheduler and job launcher combination.

3.11.3 Job schedulers and parallel launchers

```
class reframe.core.schedulers.Job(name, workdir='.', script_filename=None,
                                stdout=None, stderr=None, sched_flex_alloc_nodes=None,
                                sched_access=[], sched_account=None,
                                sched_partition=None, sched_reservation=None,
                                sched_nodelist=None, sched_exclude_nodelist=None,
                                sched_exclusive_access=None, sched_options=None)
```

Bases: `object`

A job descriptor.

A job descriptor is created by the framework after the “setup” phase and is associated with the test. It can be retrieved through the `reframe.core.pipeline.RegressionTest.job` attribute and stores information about the job submitted during the “run” phase.

Note: Users cannot create a job descriptor directly and associate it with a test.

exitcode

The exit code of the job.

This may or may not be set depending on the scheduler backend.

Type `int` or `None`.

New in version 2.21.

jobid

The ID of the current job.

Type `int` or `None`.

New in version 2.21.

launcher

The (parallel) program launcher that will be used to launch the (parallel) executable of this job.

Users are allowed to explicitly set the current job launcher, but this is only relevant in rare situations, such as when you want to wrap the current launcher command. For this specific scenario, you may have a look at the `reframe.core.launchers.LauncherWrapper` class.

The following example shows how you can replace the current partition’s launcher for this test with the “local” launcher:

```
from reframe.core.launchers.registry import getlauncher

@rfm.run_after('setup')
def set_launcher(self):
    self.job.launcher = getlauncher('local')()
```

Type `reframe.core.launchers.JobLauncher`

nodelist

The list of node names assigned to this job.

This attribute is `None` if no nodes are assigned to the job yet. This attribute is set reliably only for the `slurm` backend, i.e., Slurm *with* accounting enabled. The `squeue` scheduler backend, i.e., Slurm *without* accounting, might not set this attribute for jobs that finish very quickly. For the `local` scheduler backend, this returns an one-element list containing the hostname of the current host.

This attribute might be useful in a flexible regression test for determining the actual nodes that were assigned to the test. For more information on flexible node allocation, please refer to the corresponding [section](#) of the tutorial.

This attribute is *not* supported by the `pbs` scheduler backend.

New in version 2.17.

options

Options to be passed to the backend job scheduler.

Type `List[str]`

Default `[]`

state

The state of the job.

The value of this field is scheduler-specific.

Type `str` or `None`.

New in version 2.21.

class `reframe.core.launchers.JobLauncher`

Bases: `abc.ABC`

A job launcher.

A job launcher is the executable that actually launches a distributed program to multiple nodes, e.g., `mpirun`, `srun` etc.

Note: Users cannot create job launchers directly. You may retrieve a registered launcher backend through the `reframe.core.launchers.registry.getlauncher()` function.

Note: Changed in version 2.8: Job launchers do not get a reference to a job during their initialization.

options

List of options to be passed to the job launcher invocation.

Type `list of str`

Default `[]`

class `reframe.core.launchers.LauncherWrapper` (*target_launcher, wrapper_command, wrapper_options=[]*)

Bases: `reframe.core.launchers.JobLauncher`

Wrap a launcher object so as to modify its invocation.

This is useful for parallel debuggers. For example, to launch a regression test using the [ARM DDT](#) debugger, you can do the following:

```
@rfm.run_after('setup')
def set_launcher(self):
    self.job.launcher = LauncherWrapper(self.job.launcher, 'ddt',
                                       ['--offline'])
```

If the current system partition uses native Slurm for job submission, this setup will generate the following command in the submission script:

```
ddt --offline srun <test_executable>
```

If the current partition uses `mpirun` instead, it will generate

```
ddt --offline mpirun -np <num_tasks> ... <test_executable>
```

Parameters

- **target_launcher** – The launcher to wrap.
- **wrapper_command** – The wrapper command.
- **wrapper_options** – List of options to pass to the wrapper command.

`reframe.core.launchers.registry.getlauncher` (*name*)

Get launcher by its registered name.

The available names are those specified in the *configuration file*.

This method may become handy in very special situations, e.g., testing an application that needs to replace the system partition launcher or if a different launcher must be used for a different programming environment.

For example, if you want to replace the current partition's launcher with the local one, here is how you can achieve it:

```
def setup(self, partition, environ, **job_opts):
    super().setup(partition, environ, **job_opts)
    self.job.launcher = getlauncher('local')()
```

Note that this method returns a launcher class type and not an instance of that class. You have to instantiate it explicitly before assigning it to the `launcher` attribute of the job.

Note: New in version 2.8.

Parameters *name* – The name of the launcher to retrieve.

Returns The class of the launcher requested, which is a subclass of `reframe.core.launchers.JobLauncher`.

Raises `reframe.core.exceptions.ConfigError` – if no launcher is registered with that name.

`reframe.core.launchers.registry.register_launcher` (*name*, *local=False*)
Class decorator for registering new job launchers.

Caution: This decorator is only relevant to developers of new job launchers.

Note: New in version 2.8.

Parameters

- **name** – The registration name of this launcher
- **local** – True if launcher may only submit local jobs, False otherwise.

Raises `ValueError` – if a job launcher is already registered with the same name.

3.11.4 Runtime services

class `reframe.core.runtime.HostResources` (*prefix=None*, *stagedir=None*, *outputdir=None*,
perflogdir=None, *timefmt=None*)

Bases: `object`

Resources associated with ReFrame execution on the current host.

Note: New in version 2.13.

output_prefix

The output prefix directory of ReFrame.

prefix

The prefix directory of ReFrame execution. This is always an absolute path.

Type `str`

Caution: Users may not set this field.

stage_prefix

The stage prefix directory of ReFrame.

class `reframe.core.runtime.HostSystem` (*system*, *partname=None*)

Bases: `object`

The host system of the framework.

The host system is a representation of the system that the framework currently runs on. If the framework is properly configured, the host system is automatically detected. If not, it may be explicitly set by the user.

This class is mainly a proxy of `reframe.core.systems.System` that stores optionally a partition name and provides some additional functionality for manipulating system partitions.

All attributes of the `reframe.core.systems.System` may be accessed directly from this proxy.

Note: New in version 2.13.

partition (*name*)

Return the system partition name.

Type `reframe.core.systems.SystemPartition`.

partitions

The partitions of this system.

Type `list[reframe.core.systems.SystemPartition]`.

class `reframe.core.runtime.RuntimeContext` (*dict_config*, *sysdescr=None*, ***options*)

Bases: `object`

The runtime context of the framework.

This class essentially groups the current host system and the associated resources of the framework on the current system. It also encapsulates other runtime parameters that are relevant to the framework's execution.

There is a single instance of this class globally in the framework.

Note: New in version 2.13.

modules_system

The modules system used by the current host system.

Type `reframe.core.modules.ModulesSystem`.

non_default_craype

True if a non-default Cray PE is tested.

This will cause ReFrame to set the `LD_LIBRARY_PATH` as follows after all modules have been loaded:

```
export LD_LIBRARY_PATH=$CRAY_LD_LIBRARY_PATH:$LD_LIBRARY_PATH
```

This property is set through the `--non-default-craype` command-line option.

Type `bool` (default: `False`)

resources

The framework resources.

Type `reframe.core.runtime.HostResources`

show_config ()

Return a textual representation of the current runtime.

system

The current host system.

Type `reframe.core.runtime.HostSystem`

class `reframe.core.runtime.module_use` (**paths*)

Bases: `object`

Context manager for temporarily modifying the module path

`reframe.core.runtime.runtime` ()

Retrieve the framework's runtime context.

Type `reframe.core.runtime.RuntimeContext`

Note: New in version 2.13.

3.11.5 Modules System API

class `reframe.core.modules.ModulesSystem` (*backend*)

A modules system abstraction inside ReFrame.

This class interfaces between the framework internals and the actual modules systems implementation.

conflicted_modules (*name*)

Return the list of the modules conflicting with module *name*.

If module *name* resolves to multiple real modules, then the returned list will be the concatenation of the conflict lists of all the real modules.

This method returns a list of strings.

emit_load_commands (*name*)

Return the appropriate shell command for loading module *name*.

emit_unload_commands (*name*)

Return the appropriate shell command for unloading module *name*.

is_module_loaded (*name*)

Check if module *name* is loaded.

If module *name* refers to multiple real modules, this method will return `True` only if all the referees are loaded.

load_mapping (*mapping*)

Update the internal module mappings using a single mapping.

Parameters *mapping* – a string specifying the module mapping. Example syntax: 'm0: m1 m2'.

load_mapping_from_file (*filename*)

Update the internal module mappings from mappings read from file.

load_module (*name*, *force=False*)

Load the module *name*.

If *force* is set, forces the loading, unloading first any conflicting modules currently loaded. If module *name* refers to multiple real modules, all of the target modules will be loaded.

Returns the list of unloaded modules as strings.

loaded_modules ()

Return a list of loaded modules.

This method returns a list of strings.

name

Return the name of this module system.

resolve_module (*name*)

Resolve module *name* in the registered module map.

Returns the list of real modules names pointed to by *name*.

Raises `reframe.core.exceptions.ConfigError` if the mapping contains a cycle.

searchpath

The module system search path as a list of directories.

searchpath_add (*dirs)

Add *dirs* to the module system search path.

searchpath_remove (*dirs)

Remove *dirs* from the module system search path.

unload_all ()

Unload all loaded modules.

unload_module (name)

Unload module *name*.

If module *name* refers to multiple real modules, all the referred to modules will be unloaded in reverse order.

version

Return the version of this module system.

3.11.6 Build systems

New in version 2.14.

ReFrame delegates the compilation of the regression test to a *build system*. Build systems in ReFrame are entities that are responsible for generating the necessary shell commands for compiling a code. Each build system defines a set of attributes that users may set in order to customize their compilation. An example usage is the following:

```
self.build_system = 'SingleSource'
self.build_system.cflags = ['-fopenmp']
```

Users simply set the build system to use in their regression tests and then they configure it. If no special configuration is needed for the compilation, users may completely ignore the build systems. ReFrame will automatically pick one based on the regression test attributes and will try to compile the code.

All build systems in ReFrame derive from the abstract base class `reframe.core.buildsystems.BuildSystem`. This class defines a set of common attributes, such as compilers, compilation flags etc. that all subclasses inherit. It is up to the concrete build system implementations on how to use or not these attributes.

class reframe.core.buildsystems.Autotools

Bases: `reframe.core.buildsystems.ConfigureBasedBuildSystem`

A build system for compiling Autotools-based projects.

This build system will emit the following commands:

1. Create a build directory if `builddir` is not `None` and change to it.
2. Invoke `configure` to configure the project by setting the corresponding flags for compilers and compiler flags.
3. Issue `make` to compile the code.

emit_build_commands (environ)

Return the list of commands for building using this build system.

The build commands may always assume to be issued from the top-level directory of the code that is to be built.

Parameters `environ` (`reframe.core.environments.ProgEnvironment`) – The programming environment for which to emit the build instructions. The framework passes here the current programming environment.

Raises `BuildSystemError` in case of errors when generating the build instructions.

Note: This method is relevant only to developers of new build systems.

class `reframe.core.buildsystems.BuildSystem`

Bases: `abc.ABC`

The abstract base class of any build system.

Concrete build systems inherit from this class and must override the `emit_build_commands()` abstract function.

cc

The C compiler to be used. If set to `None` and `flags_from_environ` is `True`, the compiler defined in the current programming environment will be used.

Type `str`

Default `None`

cflags

The C compiler flags to be used. If set to `None` and `flags_from_environ` is `True`, the corresponding flags defined in the current programming environment will be used.

Type `List[str]`

Default `None`

cppflags

The preprocessor flags to be used. If set to `None` and `flags_from_environ` is `True`, the corresponding flags defined in the current programming environment will be used.

Type `List[str]`

Default `None`

cxx

The C++ compiler to be used. If set to `None` and `flags_from_environ` is `True`, the compiler defined in the current programming environment will be used.

Type `str`

Default `None`

cxxflags

The C++ compiler flags to be used. If set to `None` and `flags_from_environ` is `True`, the corresponding flags defined in the current programming environment will be used.

Type `List[str]`

Default `None`

emit_build_commands (`environ`)

Return the list of commands for building using this build system.

The build commands may always assume to be issued from the top-level directory of the code that is to be built.

Parameters `environ` (*reframe.core.environments.ProgEnvironment*) – The programming environment for which to emit the build instructions. The framework passes here the current programming environment.

Raises `BuildSystemError` in case of errors when generating the build instructions.

Note: This method is relevant only to developers of new build systems.

fflags

The Fortran compiler flags to be used. If set to `None` and `flags_from_environ` is `True`, the corresponding flags defined in the current programming environment will be used.

Type `List[str]`

Default `None`

flags_from_environ

Set compiler and compiler flags from the current programming environment if not specified otherwise.

Type `bool`

Default `True`

ftn

The Fortran compiler to be used. If set to `None` and `flags_from_environ` is `True`, the compiler defined in the current programming environment will be used.

Type `str`

Default `None`

ldflags

The linker flags to be used. If set to `None` and `flags_from_environ` is `True`, the corresponding flags defined in the current programming environment will be used.

Type `List[str]`

Default `None`

nvcc

The CUDA compiler to be used. If set to `None` and `flags_from_environ` is `True`, the compiler defined in the current programming environment will be used.

Type `str`

Default `None`

class `reframe.core.buildsystems.CMake`

Bases: *reframe.core.buildsystems.ConfigureBasedBuildSystem*

A build system for compiling CMake-based projects.

This build system will emit the following commands:

1. Create a build directory if `builddir` is not `None` and change to it.
2. Invoke `cmake` to configure the project by setting the corresponding CMake flags for compilers and compiler flags.
3. Issue `make` to compile the code.

emit_build_commands (*environ*)

Return the list of commands for building using this build system.

The build commands may always assume to be issued from the top-level directory of the code that is to be built.

Parameters `environ` (`reframe.core.environments.ProgEnvironment`) – The programming environment for which to emit the build instructions. The framework passes here the current programming environment.

Raises `BuildSystemError` in case of errors when generating the build instructions.

Note: This method is relevant only to developers of new build systems.

class `reframe.core.buildsystems.ConfigureBasedBuildSystem`

Bases: `reframe.core.buildsystems.BuildSystem`

Abstract base class for configured-based build systems.

builddir

The CMake build directory, where all the generated files will be placed.

Type `str`

Default `None`

config_opts

Additional configuration options to be passed to the CMake invocation.

Type `List[str]`

Default `[]`

make_opts

Options to be passed to the subsequent make invocation.

Type `List[str]`

Default `[]`

max_concurrency

Same as for the `Make` build system.

Type `integer`

Default `1`

srcdir

The top-level directory of the code.

This is set automatically by the framework based on the `reframe.core.pipeline.RegressionTest.sourcepath` attribute.

Type `str`

Default `None`

class `reframe.core.buildsystems.Make`

Bases: `reframe.core.buildsystems.BuildSystem`

A build system for compiling codes using make.

The generated build command has the following form:

```
make -j [N] [-f MAKEFILE] [-C SRCDIR] CC="X" CXX="X" FC="X" NVCC="X" CPPFLAGS="X" ↵
↵CFLAGS="X" CXXFLAGS="X" FCFLAGS="X" LDFLAGS="X" OPTIONS
```

The compiler and compiler flags variables will only be passed if they are not `None`. Their value is determined by the corresponding attributes of `BuildSystem`. If you want to completely disable passing these variables to the `make` invocation, you should make sure not to set any of the corresponding attributes and set also the `BuildSystem.flags_from_extern` flag to `False`.

emit_build_commands (*environ*)

Return the list of commands for building using this build system.

The build commands may always assume to be issued from the top-level directory of the code that is to be built.

Parameters `environ` (`reframe.core.environments.ProgEnvironment`) – The programming environment for which to emit the build instructions. The framework passes here the current programming environment.

Raises `BuildSystemError` in case of errors when generating the build instructions.

Note: This method is relevant only to developers of new build systems.

makefile

Instruct build system to use this Makefile. This option is useful when having non-standard Makefile names.

Type `str`

Default `None`

max_concurrency

Limit concurrency for `make` jobs. This attribute controls the `-j` option passed to `make`. If not `None`, `make` will be invoked as `make -j max_concurrency`. Otherwise, it will invoked as `make -j`.

Type `integer`

Default `1`

Note: Changed in version 2.19: The default value is now `1`

options

Append these options to the `make` invocation. This variable is also useful for passing variables or targets to `make`.

Type `List[str]`

Default `[]`

srcdir

The top-level directory of the code.

This is set automatically by the framework based on the `reframe.core.pipeline.RegressionTest.sourcepath` attribute.

Type `str`

Default `None`

class `reframe.core.buildsystems.SingleSource`

Bases: `reframe.core.buildsystems.BuildSystem`

A build system for compiling a single source file.

The generated build command will have the following form:

```
COMP CPPFLAGS XFLAGS SRCFILE -o EXEC LDFLAGS
```

- `COMP` is the required compiler for compiling `SRCFILE`. This build system will automatically detect the programming language of the source file and pick the correct compiler. See also the *SingleSource.lang* attribute.
- `CPPFLAGS` are the preprocessor flags and are passed to any compiler.
- `XFLAGS` is any of `CFLAGS`, `CXXFLAGS` or `FCFLAGS` depending on the programming language of the source file.
- `SRCFILE` is the source file to be compiled. This is set up automatically by the framework. See also the *SingleSource.srcfile* attribute.
- `EXEC` is the executable to be generated. This is also set automatically by the framework. See also the *SingleSource.executable* attribute.
- `LDFLAGS` are the linker flags.

For CUDA codes, the language assumed is C++ (for the compilation flags) and the compiler used is *BuildSystem.nvcc*.

emit_build_commands (*environ*)

Return the list of commands for building using this build system.

The build commands may always assume to be issued from the top-level directory of the code that is to be built.

Parameters *environ* (*reframe.core.environments.ProgEnvironment*) – The programming environment for which to emit the build instructions. The framework passes here the current programming environment.

Raises `BuildSystemError` in case of errors when generating the build instructions.

Note: This method is relevant only to developers of new build systems.

executable

The executable file to be generated.

This is set automatically by the framework based on the *reframe.core.pipeline.RegressionTest.executable* attribute.

Type `str` or `None`

include_path

The include path to be used for this compilation.

All the elements of this list will be appended to the *BuildSystem.cppflags*, by prepending to each of them the `-I` option.

Type `List[str]`

Default `[]`

lang

The programming language of the file that needs to be compiled. If not specified, the build system will try to figure it out automatically based on the extension of the source file. The automatically detected extensions are the following:

- C: `.c`.

- C++: `.cc`, `.cp`, `.cxx`, `.cpp`, `.CPP`, `.c++` and `.C`.
- Fortran: `.f`, `.for`, `.ftn`, `.F`, `.FOR`, `.fpp`, `.FPP`, `.FTN`, `.f90`, `.f95`, `.f03`, `.f08`, `.F90`, `.F95`, `.F03` and `.F08`.
- CUDA: `.cu`.

Type `str` or `None`

srcfile

The source file to compile. This is automatically set by the framework based on the `reframe.core.pipeline.RegressionTest.sourcepath` attribute.

Type `str` or `None`

3.11.7 Container platforms

New in version 2.20.

ReFrame can run a regression test inside a container. To achieve that you have to set the `reframe.core.pipeline.RegressionTest.container_platform` attribute and then set up the container platform (e.g., image to load, commands to execute). The `reframe.core.ContainerPlatform` abstract base class define the basic interface and a minimal set of attributes that all concrete container platforms must implement. Concrete container platforms may also define additional fields that are specific to them.

class `reframe.core.containers.ContainerPlatform`

Bases: `abc.ABC`

The abstract base class of any container platform.

Concrete container platforms inherit from this class and must override the `emit_prepare_commands()` and `launch_command()` abstract methods.

commands

The commands to be executed within the container.

Type `list[str]`

Default `[]`

emit_prepare_commands()

Returns commands for preparing this container for running.

Such a command could be for pulling the container image from a repository.

image

The container image to be used for running the test.

Type `str` or `None`

Default `None`

launch_command()

Returns the command for running `commands` with this container platform.

mount_points

List of mount point pairs for directories to mount inside the container.

Each mount point is specified as a tuple of `(/path/in/host, /path/in/container)`.

Type `list[tuple[str, str]]`

Default `[]`

options

Additional options to be passed to the container runtime when executed.

Type `list[str]`

Default `[]`

workdir

The working directory of ReFrame inside the container.

This is the directory where the test's stage directory is mounted inside the container. This directory is always mounted regardless if `mount_points` is set or not.

Type `str`

Default `/rfm_workdir`

class `reframe.core.containers.Docker`

Bases: `reframe.core.containers.ContainerPlatform`

Container platform backend for running containers with `Docker`.

emit_prepare_commands ()

Returns commands for preparing this container for running.

Such a command could be for pulling the container image from a repository.

launch_command ()

Returns the command for running `commands` with this container platform.

class `reframe.core.containers.Sarus`

Bases: `reframe.core.containers.ContainerPlatform`

Container platform backend for running containers with `Sarus`.

emit_prepare_commands ()

Returns commands for preparing this container for running.

Such a command could be for pulling the container image from a repository.

launch_command ()

Returns the command for running `commands` with this container platform.

with_mpi

Enable MPI support when launching the container.

Type `boolean`

Default `False`

class `reframe.core.containers.ShifterNG`

Bases: `reframe.core.containers.Sarus`

Container platform backend for running containers with `ShifterNG`.

class `reframe.core.containers.Singularity`

Bases: `reframe.core.containers.ContainerPlatform`

Container platform backend for running containers with `Singularity`.

emit_prepare_commands ()

Returns commands for preparing this container for running.

Such a command could be for pulling the container image from a repository.

launch_command ()

Returns the command for running `commands` with this container platform.

with_cuda

Enable CUDA support when launching the container.

Type boolean

Default False

3.12 Sanity Functions Reference

Sanity deferrable functions.

This module provides functions to be used with the `sanity_patterns` and `:attr'perf_patterns <reframe.core.pipeline.RegressionTest.perf_patterns>'`. The key characteristic of these functions is that they are not executed the time they are called. Instead they are evaluated at a later point by the framework (inside the `check_sanity` and `check_performance` methods). Any sanity function may be evaluated either explicitly or implicitly.

3.12.1 Explicit evaluation of sanity functions

Sanity functions may be evaluated at any time by calling the `evaluate` on their return value.

3.12.2 Implicit evaluation of sanity functions

Sanity functions may also be evaluated implicitly in the following situations:

- When you try to get their truthy value by either explicitly or implicitly calling `bool` on their return value. This implies that when you include the result of a sanity function in an `if` statement or when you apply the `and`, `or` or `not` operators, this will trigger their immediate evaluation.
- When you try to iterate over their result. This implies that including the result of a sanity function in a `for` statement will trigger its evaluation immediately.
- When you try to explicitly or implicitly get its string representation by calling `str` on its result. This implies that printing the return value of a sanity function will automatically trigger its evaluation.

This module provides three categories of sanity functions:

1. Deferrable replacements of certain Python built-in functions. These functions simply delegate their execution to the actual built-ins.
2. Assertion functions. These functions are used to assert certain conditions and they either return `True` or raise `reframe.core.exceptions.SanityError` with a message describing the error. Users may provide their own formatted messages through the `msg` argument. For example, in the following call to `assert_eq()` the `{0}` and `{1}` placeholders will obtain the actual arguments passed to the assertion function.

```
assert_eq(a, 1, msg="{0} is not equal to {1}")
```

If in the user provided message more placeholders are used than the arguments of the assert function (except the `msg` argument), no argument substitution will be performed in the user message.

3. Utility functions. These are functions that you will normally use when defining `sanity_patterns` and `perf_patterns`. They include, but are not limited to, functions to iterate over regex matches in a file, extracting and converting values from regex matches, computing statistical information on series of data etc.

`reframe.utility.sanity.abs(x)`

Replacement for the built-in `abs()` function.

`reframe.utility.sanity.all(iterable)`

Replacement for the built-in `all()` function.

`reframe.utility.sanity.allx(iterable)`

Same as the built-in `all()` function, except that it returns `False` if `iterable` is empty.

New in version 2.13.

`reframe.utility.sanity.and_(a, b)`

Deferrable version of the `and` operator.

Returns `a` and `b`.

`reframe.utility.sanity.any(iterable)`

Replacement for the built-in `any()` function.

`reframe.utility.sanity.assert_bounded(val, lower=None, upper=None, msg=None)`

Assert that `lower <= val <= upper`.

Parameters

- **val** – The value to check.
- **lower** – The lower bound. If `None`, it defaults to `-inf`.
- **upper** – The upper bound. If `None`, it defaults to `inf`.

Returns `True` on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

`reframe.utility.sanity.assert_eq(a, b, msg=None)`

Assert that `a == b`.

Returns `True` on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

`reframe.utility.sanity.assert_false(x, msg=None)`

Assert that `x` is evaluated to `False`.

Returns `True` on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

`reframe.utility.sanity.assert_found(patt, filename, msg=None, encoding='utf-8')`

Assert that regex pattern `patt` is found in the file `filename`.

Parameters

- **patt** – The regex pattern to search. Any standard Python [regular expression](#) is accepted.
- **filename** – The name of the file to examine. Any `OSError` raised while processing the file will be propagated as a `reframe.core.exceptions.SanityError`.
- **encoding** – The name of the encoding used to decode the file.

Returns `True` on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

`reframe.utility.sanity.assert_ge(a, b, msg=None)`

Assert that `a >= b`.

Returns `True` on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

```
reframe.utility.sanity.assert_gt(a, b, msg=None)
```

Assert that $a > b$.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

```
reframe.utility.sanity.assert_in(item, container, msg=None)
```

Assert that `item` is in `container`.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

```
reframe.utility.sanity.assert_le(a, b, msg=None)
```

Assert that $a \leq b$.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

```
reframe.utility.sanity.assert_lt(a, b, msg=None)
```

Assert that $a < b$.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

```
reframe.utility.sanity.assert_ne(a, b, msg=None)
```

Assert that $a \neq b$.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

```
reframe.utility.sanity.assert_not_found(patt, filename, msg=None, encoding='utf-8')
```

Assert that regex pattern `patt` is not found in the file `filename`.

This is the inverse of `assert_found()`.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

```
reframe.utility.sanity.assert_not_in(item, container, msg=None)
```

Assert that `item` is not in `container`.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

```
reframe.utility.sanity.assert_reference(val, ref, lower_thres=None, upper_thres=None,
                                       msg=None)
```

Assert that value `val` respects the reference value `ref`.

Parameters

- **val** – The value to check.
- **ref** – The reference value.
- **lower_thres** – The lower threshold value expressed as a negative decimal fraction of the reference value. Must be in $[-1, 0]$ for $ref \geq 0.0$ and in $[-inf, 0]$ for $ref < 0.0$. If `None`, no lower thresholds is applied.
- **upper_thres** – The upper threshold value expressed as a decimal fraction of the reference value. Must be in $[0, inf]$ for $ref \geq 0.0$ and in $[0, 1]$ for $ref < 0.0$. If `None`, no upper thresholds is applied.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails or if the lower and upper thresholds do not have appropriate values.

`reframe.utility.sanity.assert_true(x, msg=None)`
Assert that `x` is evaluated to True.

Returns True on success.

Raises `reframe.core.exceptions.SanityError` – if assertion fails.

`reframe.utility.sanity.avg(iterable)`
Return the average of all the elements of `iterable`.

`reframe.utility.sanity.chain(*iterables)`
Replacement for the `itertools.chain()` function.

`reframe.utility.sanity.contains(seq, key)`
Deferrable version of the `in` operator.

Returns `key in seq`.

`reframe.utility.sanity.count(iterable)`
Return the element count of `iterable`.

This is similar to the built-in `len()`, except that it can also handle any argument that supports iteration, including generators.

`reframe.utility.sanity.count_uniq(iterable)`
Return the unique element count of `iterable`.

`reframe.utility.sanity.defer(x)`
Defer the evaluation of variable `x`.

New in version 2.21.

`reframe.utility.sanity.enumerate(iterable, start=0)`
Replacement for the built-in `enumerate()` function.

`reframe.utility.sanity.evaluate(expr)`
Evaluate a deferred expression.

If `expr` is not a deferred expression, it will be returned as is.

New in version 2.21.

`reframe.utility.sanity.extractall(patt, filename, tag=0, conv=None, encoding='utf-8')`
Extract all values from the capturing group `tag` of a matching regex `patt` in the file `filename`.

Parameters

- **patt** – The regex pattern to search. Any standard Python [regular expression](#) is accepted.
- **filename** – The name of the file to examine.
- **encoding** – The name of the encoding used to decode the file.
- **tag** – The regex capturing group to be extracted. Group 0 refers always to the whole match. Since the file is processed line by line, this means that group 0 returns the whole line that was matched.
- **conv** – A callable that takes a single argument and returns a new value. If provided, it will be used to convert the extracted values before returning them.

Returns A list of the extracted values from the matched regex.

Raises `reframe.core.exceptions.SanityError` – In case of errors.

```
reframe.utility.sanity.extractiter (patt, filename, tag=0, conv=None, encoding='utf-8')
```

Get an iterator over the values extracted from the capturing group `tag` of a matching regex `patt` in the file `filename`.

This function is equivalent to `extractall()` except that it returns a generator object, instead of a list, which you can use to iterate over the extracted values.

```
reframe.utility.sanity.extractsingle (patt, filename, tag=0, conv=None, item=0,
                                       encoding='utf-8')
```

Extract a single value from the capturing group `tag` of a matching regex `patt` in the file `filename`.

This function is equivalent to `extractall(patt, filename, tag, conv)[item]`, except that it raises a `SanityError` if `item` is out of bounds.

Parameters

- **patt** – as in `extractall()`.
- **filename** – as in `extractall()`.
- **encoding** – as in `extractall()`.
- **tag** – as in `extractall()`.
- **conv** – as in `extractall()`.
- **item** – the specific element to extract.

Returns The extracted value.

Raises `reframe.core.exceptions.SanityError` – In case of errors.

```
reframe.utility.sanity.filter (function, iterable)
```

Replacement for the built-in `filter()` function.

```
reframe.utility.sanity.findall (patt, filename, encoding='utf-8')
```

Get all matches of regex `patt` in `filename`.

Parameters

- **patt** – The regex pattern to search. Any standard Python [regular expression](#) is accepted.
- **filename** – The name of the file to examine.
- **encoding** – The name of the encoding used to decode the file.

Returns A list of raw [regex match objects](#).

Raises `reframe.core.exceptions.SanityError` – In case an `OSError` is raised while processing `filename`.

```
reframe.utility.sanity.finditer (patt, filename, encoding='utf-8')
```

Get an iterator over the matches of the regex `patt` in `filename`.

This function is equivalent to `findall()` except that it returns a generator object instead of a list, which you can use to iterate over the raw matches.

```
reframe.utility.sanity.getattr (obj, attr, *args)
```

Replacement for the built-in `getattr()` function.

```
reframe.utility.sanity.getitem (container, item)
```

Get `item` from `container`.

`container` may refer to any container that can be indexed.

Raises `reframe.core.exceptions.SanityError` – In case `item` cannot be retrieved from container.

`reframe.utility.sanity.glob(pathname, *, recursive=False)`

Replacement for the `glob.glob()` function.

`reframe.utility.sanity.hasattr(obj, name)`

Replacement for the built-in `hasattr()` function.

`reframe.utility.sanity.iglob(pathname, recursive=False)`

Replacement for the `glob.iglob()` function.

`reframe.utility.sanity.len(s)`

Replacement for the built-in `len()` function.

`reframe.utility.sanity.map(function, *iterables)`

Replacement for the built-in `map()` function.

`reframe.utility.sanity.max(*args)`

Replacement for the built-in `max()` function.

`reframe.utility.sanity.min(*args)`

Replacement for the built-in `min()` function.

`reframe.utility.sanity.not_(a)`

Deferrable version of the `not` operator.

Returns `not a`.

`reframe.utility.sanity.or_(a, b)`

Deferrable version of the `or` operator.

Returns `a or b`.

`reframe.utility.sanity.print(*objects, sep=' ', end='\n', file=None, flush=False)`

Replacement for the built-in `print()` function.

The only difference is that this function returns the `objects`, so that you can use it transparently inside a complex sanity expression. For example, you could write the following to print the matches returned from the `extractall()` function:

```
self.sanity_patterns = sn.assert_eq(
    sn.count(sn.print(sn.extract_all(...))), 10
)
```

If `file` is `None`, `print()` will print its arguments to the standard output. Unlike the builtin `print()` function, we don't bind the `file` argument to `sys.stdout` by default. This would capture `sys.stdout` at the time this function is defined and would prevent it from seeing changes to `sys.stdout`, such as redirects, in the future.

`reframe.utility.sanity.reversed(seq)`

Replacement for the built-in `reversed()` function.

`reframe.utility.sanity.round(number, *args)`

Replacement for the built-in `round()` function.

`reframe.utility.sanity.sanity_function(func)`

Decorator Sanity function decorator.

Decorate any function to be used in sanity and/or performance patterns with this decorator:

```
@sanity_function
def myfunc(*args):
    do_sth()
```

This decorator is an alias to the `reframe.core.deferrable.deferrable()` decorator. The following function definition is equivalent to the above:

```
@deferrable
def myfunc(*args):
    do_sth()
```

`reframe.utility.sanity.setattr(obj, name, value)`

Replacement for the built-in `setattr()` function.

`reframe.utility.sanity.sorted(iterable, *args)`

Replacement for the built-in `sorted()` function.

`reframe.utility.sanity.sum(iterable, *args)`

Replacement for the built-in `sum()` function.

`reframe.utility.sanity.zip(*iterables)`

Replacement for the built-in `zip()` function.

r

- `reframe.core.buildsystems`, 117
- `reframe.core.containers`, 123
- `reframe.core.decorators`, 95
- `reframe.core.environments`, 108
- `reframe.core.launchers`, 112
- `reframe.core.launchers.registry`, 113
- `reframe.core.pipeline`, 96
- `reframe.core.runtime`, 114
- `reframe.core.systems`, 109
- `reframe.utility.sanity`, 125

A

abs() (in module *reframe.utility.sanity*), 125
 all() (in module *reframe.utility.sanity*), 125
 allx() (in module *reframe.utility.sanity*), 126
 and_() (in module *reframe.utility.sanity*), 126
 any() (in module *reframe.utility.sanity*), 126
 assert_bounded() (in module *reframe.utility.sanity*), 126
 assert_eq() (in module *reframe.utility.sanity*), 126
 assert_false() (in module *reframe.utility.sanity*), 126
 assert_found() (in module *reframe.utility.sanity*), 126
 assert_ge() (in module *reframe.utility.sanity*), 126
 assert_gt() (in module *reframe.utility.sanity*), 126
 assert_in() (in module *reframe.utility.sanity*), 127
 assert_le() (in module *reframe.utility.sanity*), 127
 assert_lt() (in module *reframe.utility.sanity*), 127
 assert_ne() (in module *reframe.utility.sanity*), 127
 assert_not_found() (in module *reframe.utility.sanity*), 127
 assert_not_in() (in module *reframe.utility.sanity*), 127
 assert_reference() (in module *reframe.utility.sanity*), 127
 assert_true() (in module *reframe.utility.sanity*), 128
 Autotools (class in *reframe.core.buildsystems*), 117
 avg() (in module *reframe.utility.sanity*), 128

B

build_system (*reframe.core.pipeline.RegressionTest* attribute), 96
 builddir (*reframe.core.buildsystems.ConfigureBasedBuildSystem* attribute), 120
 BuildSystem (class in *reframe.core.buildsystems*), 118

C

cc (*reframe.core.buildsystems.BuildSystem* attribute), 118

cc (*reframe.core.environments.ProgEnvironment* attribute), 108
 cflags (*reframe.core.buildsystems.BuildSystem* attribute), 118
 cflags (*reframe.core.environments.ProgEnvironment* attribute), 109
 chain() (in module *reframe.utility.sanity*), 128
 check_performance() (*reframe.core.pipeline.RegressionTest* method), 97
 check_sanity() (*reframe.core.pipeline.RegressionTest* method), 97
 cleanup() (*reframe.core.pipeline.RegressionTest* method), 97
 CMake (class in *reframe.core.buildsystems*), 119
 commands (*reframe.core.containers.ContainerPlatform* attribute), 123
 compile() (*reframe.core.pipeline.RegressionTest* method), 97
 compile() (*reframe.core.pipeline.RunOnlyRegressionTest* method), 107
 compile_wait() (*reframe.core.pipeline.RegressionTest* method), 97
 compile_wait() (*reframe.core.pipeline.RunOnlyRegressionTest* method), 107
 CompileOnlyRegressionTest (class in *reframe.core.pipeline*), 107
 config_opts (*reframe.core.buildsystems.ConfigureBasedBuildSystem* attribute), 120
 ConfigureBasedBuildSystem (class in *reframe.core.buildsystems*), 120
 conflicted_modules() (*reframe.core.modules.ModulesSystem* method), 116
 container_platform (*reframe.core.pipeline.RegressionTest* attribute), 97

- ContainerPlatform (class in reframe.core.containers), 123
- contains() (in module reframe.utility.sanity), 128
- count() (in module reframe.utility.sanity), 128
- count_uniq() (in module reframe.utility.sanity), 128
- cppflags (reframe.core.buildsystems.BuildSystem attribute), 118
- cppflags (reframe.core.environments.ProgEnvironment attribute), 109
- current_envIRON (reframe.core.pipeline.RegressionTest attribute), 97
- current_partition (reframe.core.pipeline.RegressionTest attribute), 98
- current_system (reframe.core.pipeline.RegressionTest attribute), 98
- cxx (reframe.core.buildsystems.BuildSystem attribute), 118
- cxx (reframe.core.environments.ProgEnvironment attribute), 109
- cxxflags (reframe.core.buildsystems.BuildSystem attribute), 118
- cxxflags (reframe.core.environments.ProgEnvironment attribute), 109
- ## D
- defer() (in module reframe.utility.sanity), 128
- DEPEND_BY_ENV (in module reframe.core.pipeline), 108
- DEPEND_EXACT (in module reframe.core.pipeline), 107
- DEPEND_FULLY (in module reframe.core.pipeline), 108
- depends_on() (reframe.core.pipeline.RegressionTest method), 98
- descr (reframe.core.pipeline.RegressionTest attribute), 98
- descr (reframe.core.systems.System attribute), 109
- descr (reframe.core.systems.SystemPartition attribute), 110
- details() (reframe.core.environments.Environment method), 108
- details() (reframe.core.environments.ProgEnvironment method), 109
- Docker (class in reframe.core.containers), 124
- ## E
- emit_build_commands() (reframe.core.buildsystems.Autotools method), 117
- emit_build_commands() (reframe.core.buildsystems.BuildSystem method), 118
- emit_build_commands() (reframe.core.buildsystems.CMake method), 119
- emit_build_commands() (reframe.core.buildsystems.Make method), 121
- emit_build_commands() (reframe.core.buildsystems.SingleSource method), 122
- emit_load_commands() (reframe.core.modules.ModulesSystem method), 116
- emit_prepare_commands() (reframe.core.containers.ContainerPlatform method), 123
- emit_prepare_commands() (reframe.core.containers.Docker method), 124
- emit_prepare_commands() (reframe.core.containers.Sarus method), 124
- emit_prepare_commands() (reframe.core.containers.Singularity method), 124
- emit_unload_commands() (reframe.core.modules.ModulesSystem method), 116
- enumerate() (in module reframe.utility.sanity), 128
- Environment (class in reframe.core.environments), 108
- evaluate() (in module reframe.utility.sanity), 128
- exclusive_access (reframe.core.pipeline.RegressionTest attribute), 98
- executable (reframe.core.buildsystems.SingleSource attribute), 122
- executable (reframe.core.pipeline.RegressionTest attribute), 98
- executable_opts (reframe.core.pipeline.RegressionTest attribute), 98
- exitcode (reframe.core.schedulers.Job attribute), 111
- extra_resources (reframe.core.pipeline.RegressionTest attribute), 99
- extractall() (in module reframe.utility.sanity), 128
- extractiter() (in module reframe.utility.sanity), 129
- extractsingle() (in module reframe.utility.sanity), 129
- ## F
- fflags (reframe.core.buildsystems.BuildSystem attribute), 119
- fflags (reframe.core.environments.ProgEnvironment attribute), 109
- filter() (in module reframe.utility.sanity), 129

findall() (in module *reframe.utility.sanity*), 129
 finditer() (in module *reframe.utility.sanity*), 129
 flags_from_environ (re-
frame.core.buildsystems.BuildSystem at-
 tribute), 119
 ftn (*reframe.core.buildsystems.BuildSystem* attribute),
 119
 ftn (*reframe.core.environments.ProgEnvironment*
 attribute), 109
 fullname (*reframe.core.systems.SystemPartition*
 attribute), 110

G

getattr() (in module *reframe.utility.sanity*), 129
 getdep() (*reframe.core.pipeline.RegressionTest*
 method), 99
 getitem() (in module *reframe.utility.sanity*), 129
 getlauncher() (in module re-
frame.core.launchers.registry), 113
 glob() (in module *reframe.utility.sanity*), 130

H

hasattr() (in module *reframe.utility.sanity*), 130
 hostnames (*reframe.core.systems.System* attribute),
 110
 HostResources (class in *reframe.core.runtime*), 114
 HostSystem (class in *reframe.core.runtime*), 114

I

iglob() (in module *reframe.utility.sanity*), 130
 image (*reframe.core.containers.ContainerPlatform* at-
 tribute), 123
 include_path (*reframe.core.buildsystems.SingleSource*
 attribute), 122
 info() (*reframe.core.pipeline.RegressionTest* method),
 100
 is_loaded (*reframe.core.environments.Environment*
 attribute), 108
 is_local() (*reframe.core.pipeline.RegressionTest*
 method), 100
 is_module_loaded() (re-
frame.core.modules.ModulesSystem method),
 116

J

Job (class in *reframe.core.schedulers*), 111
 job (*reframe.core.pipeline.RegressionTest* attribute), 100
 jobid (*reframe.core.schedulers.Job* attribute), 111
 JobLauncher (class in *reframe.core.launchers*), 112

K

keep_files (*reframe.core.pipeline.RegressionTest* at-
 tribute), 100

L

lang (*reframe.core.buildsystems.SingleSource* at-
 tribute), 122
 launch_command() (re-
frame.core.containers.ContainerPlatform
 method), 123
 launch_command() (*reframe.core.containers.Docker*
 method), 124
 launch_command() (*reframe.core.containers.Sarus*
 method), 124
 launch_command() (re-
frame.core.containers.Singularity method),
 124
 launcher (*reframe.core.schedulers.Job* attribute), 111
 launcher (*reframe.core.systems.SystemPartition*
 attribute), 110
 LauncherWrapper (class in *reframe.core.launchers*),
 113
 ldflags (*reframe.core.buildsystems.BuildSystem* at-
 tribute), 119
 ldflags (*reframe.core.environments.ProgEnvironment*
 attribute), 109
 len() (in module *reframe.utility.sanity*), 130
 load() (in module *reframe.core.environments*), 109
 load_mapping() (re-
frame.core.modules.ModulesSystem method),
 116
 load_mapping_from_file() (re-
frame.core.modules.ModulesSystem method),
 116
 load_module() (re-
frame.core.modules.ModulesSystem method),
 116
 loaded_modules() (re-
frame.core.modules.ModulesSystem method),
 116
 local (*reframe.core.pipeline.RegressionTest* attribute),
 100
 logger (*reframe.core.pipeline.RegressionTest* at-
 tribute), 101

M

maintainers (*reframe.core.pipeline.RegressionTest*
 attribute), 101
 Make (class in *reframe.core.buildsystems*), 120
 make_opts (*reframe.core.buildsystems.ConfigureBasedBuildSystem*
 attribute), 120
 makefile (*reframe.core.buildsystems.Make* attribute),
 121
 map() (in module *reframe.utility.sanity*), 130
 max() (in module *reframe.utility.sanity*), 130
 max_concurrency (re-
frame.core.buildsystems.ConfigureBasedBuildSystem
 attribute), 120

- max_concurrency (*reframe.core.buildsystems.Make attribute*), 121
 min() (*in module reframe.utility.sanity*), 130
 module_use (*class in reframe.core.runtime*), 115
 modules (*reframe.core.environments.Environment attribute*), 108
 modules (*reframe.core.pipeline.RegressionTest attribute*), 101
 modules_system (*reframe.core.runtime.RuntimeContext attribute*), 115
 modules_system (*reframe.core.systems.System attribute*), 110
 ModulesSystem (*class in reframe.core.modules*), 116
 mount_points (*reframe.core.containers.ContainerPlatform attribute*), 123
- ## N
- name (*reframe.core.environments.Environment attribute*), 108
 name (*reframe.core.modules.ModulesSystem attribute*), 116
 name (*reframe.core.pipeline.RegressionTest attribute*), 101
 name (*reframe.core.systems.System attribute*), 110
 name (*reframe.core.systems.SystemPartition attribute*), 111
 nodelist (*reframe.core.schedulers.Job attribute*), 112
 non_default_craype (*reframe.core.runtime.RuntimeContext attribute*), 115
 not() (*in module reframe.utility.sanity*), 130
 num_cpus_per_task (*reframe.core.pipeline.RegressionTest attribute*), 101
 num_gpus_per_node (*reframe.core.pipeline.RegressionTest attribute*), 101
 num_tasks (*reframe.core.pipeline.RegressionTest attribute*), 101
 num_tasks_per_core (*reframe.core.pipeline.RegressionTest attribute*), 102
 num_tasks_per_node (*reframe.core.pipeline.RegressionTest attribute*), 102
 num_tasks_per_socket (*reframe.core.pipeline.RegressionTest attribute*), 102
 nvcc (*reframe.core.buildsystems.BuildSystem attribute*), 119
- ## O
- options (*reframe.core.buildsystems.Make attribute*), 121
 options (*reframe.core.containers.ContainerPlatform attribute*), 123
 options (*reframe.core.launchers.JobLauncher attribute*), 112
 options (*reframe.core.schedulers.Job attribute*), 112
 or() (*in module reframe.utility.sanity*), 130
 output_prefix (*reframe.core.runtime.HostResources attribute*), 114
 outputdir (*reframe.core.pipeline.RegressionTest attribute*), 102
 outputdir (*reframe.core.systems.System attribute*), 110
- ## P
- parameterized_test() (*in module reframe.core.decorators*), 95
 partition() (*reframe.core.runtime.HostSystem method*), 115
 partitions (*reframe.core.runtime.HostSystem attribute*), 115
 partitions (*reframe.core.systems.System attribute*), 110
 perf_patterns (*reframe.core.pipeline.RegressionTest attribute*), 102
 perflogdir (*reframe.core.systems.System attribute*), 110
 poll() (*reframe.core.pipeline.RegressionTest method*), 102
 post_run (*reframe.core.pipeline.RegressionTest attribute*), 102
 postbuild_cmd (*reframe.core.pipeline.RegressionTest attribute*), 103
 pre_run (*reframe.core.pipeline.RegressionTest attribute*), 103
 prebuild_cmd (*reframe.core.pipeline.RegressionTest attribute*), 103
 prefix (*reframe.core.pipeline.RegressionTest attribute*), 103
 prefix (*reframe.core.runtime.HostResources attribute*), 114
 prefix (*reframe.core.systems.System attribute*), 110
 preload_envron (*reframe.core.systems.System attribute*), 110
 print() (*in module reframe.utility.sanity*), 130
 ProgEnvironment (*class in reframe.core.environments*), 108
- ## R
- readonly_files (*reframe.core.pipeline.RegressionTest attribute*), 103

- reference (*reframe.core.pipeline.RegressionTest* attribute), 103
- reframe.core.buildsystems* (module), 117
- reframe.core.containers* (module), 123
- reframe.core.decorators* (module), 95
- reframe.core.environments* (module), 108
- reframe.core.launchers* (module), 112
- reframe.core.launchers.registry* (module), 113
- reframe.core.pipeline* (module), 96
- reframe.core.runtime* (module), 114
- reframe.core.systems* (module), 109
- reframe.utility.sanity* (module), 125
- register_launcher() (in module *reframe.core.launchers.registry*), 114
- RegressionTest (class in *reframe.core.pipeline*), 96
- require_deps() (in module *reframe.core.decorators*), 96
- required_version() (in module *reframe.core.decorators*), 95
- resolve_module() (*reframe.core.modules.ModulesSystem* method), 116
- resources (*reframe.core.runtime.RuntimeContext* attribute), 115
- resourcesdir (*reframe.core.systems.System* attribute), 110
- reversed() (in module *reframe.utility.sanity*), 130
- round() (in module *reframe.utility.sanity*), 130
- run() (*reframe.core.pipeline.CompileOnlyRegressionTest* method), 107
- run() (*reframe.core.pipeline.RegressionTest* method), 104
- run() (*reframe.core.pipeline.RunOnlyRegressionTest* method), 107
- run_after() (in module *reframe.core.decorators*), 96
- run_before() (in module *reframe.core.decorators*), 96
- RunOnlyRegressionTest (class in *reframe.core.pipeline*), 106
- runtime() (in module *reframe.core.runtime*), 115
- RuntimeContext (class in *reframe.core.runtime*), 115
- ## S
- sanity_function() (in module *reframe.utility.sanity*), 130
- sanity_patterns (*reframe.core.pipeline.RegressionTest* attribute), 104
- Sarus (class in *reframe.core.containers*), 124
- scheduler (*reframe.core.systems.SystemPartition* attribute), 111
- searchpath (*reframe.core.modules.ModulesSystem* attribute), 116
- searchpath_add() (*reframe.core.modules.ModulesSystem* method), 117
- searchpath_remove() (*reframe.core.modules.ModulesSystem* method), 117
- setattr() (in module *reframe.utility.sanity*), 131
- setup() (*reframe.core.pipeline.CompileOnlyRegressionTest* method), 107
- setup() (*reframe.core.pipeline.RegressionTest* method), 104
- ShifterNG (class in *reframe.core.containers*), 124
- show_config() (*reframe.core.runtime.RuntimeContext* method), 115
- simple_test() (in module *reframe.core.decorators*), 95
- SingleSource (class in *reframe.core.buildsystems*), 121
- Singularity (class in *reframe.core.containers*), 124
- snapshot() (in module *reframe.core.environments*), 109
- sorted() (in module *reframe.utility.sanity*), 131
- sourcepath (*reframe.core.pipeline.RegressionTest* attribute), 104
- sourcesdir (*reframe.core.pipeline.RegressionTest* attribute), 105
- srcdir (*reframe.core.buildsystems.ConfigureBasedBuildSystem* attribute), 120
- srcdir (*reframe.core.buildsystems.Make* attribute), 121
- srcfile (*reframe.core.buildsystems.SingleSource* attribute), 123
- stage_prefix (*reframe.core.runtime.HostResources* attribute), 114
- stagedir (*reframe.core.pipeline.RegressionTest* attribute), 105
- stagedir (*reframe.core.systems.System* attribute), 110
- state (*reframe.core.schedulers.Job* attribute), 112
- stderr (*reframe.core.pipeline.CompileOnlyRegressionTest* attribute), 107
- stderr (*reframe.core.pipeline.RegressionTest* attribute), 105
- stdout (*reframe.core.pipeline.CompileOnlyRegressionTest* attribute), 107
- stdout (*reframe.core.pipeline.RegressionTest* attribute), 105
- strict_check (*reframe.core.pipeline.RegressionTest* attribute), 105
- sum() (in module *reframe.utility.sanity*), 131
- System (class in *reframe.core.systems*), 109
- system (*reframe.core.runtime.RuntimeContext* attribute), 115
- SystemPartition (class in *reframe.core.systems*), 110

T

`tags` (*reframe.core.pipeline.RegressionTest* attribute), 105

`temp_environment` (class in *reframe.core.environments*), 109

`time_limit` (*reframe.core.pipeline.RegressionTest* attribute), 106

U

`unload_all()` (*reframe.core.modules.ModulesSystem* method), 117

`unload_module()` (*reframe.core.modules.ModulesSystem* method), 117

`use_multithreading` (*reframe.core.pipeline.RegressionTest* attribute), 106

V

`valid_prog_environs` (*reframe.core.pipeline.RegressionTest* attribute), 106

`valid_systems` (*reframe.core.pipeline.RegressionTest* attribute), 106

`variables` (*reframe.core.environments.Environment* attribute), 108

`variables` (*reframe.core.pipeline.RegressionTest* attribute), 106

`version` (*reframe.core.modules.ModulesSystem* attribute), 117

W

`wait()` (*reframe.core.pipeline.CompileOnlyRegressionTest* method), 107

`wait()` (*reframe.core.pipeline.RegressionTest* method), 106

`with_cuda` (*reframe.core.containers.Singularity* attribute), 124

`with_mpi` (*reframe.core.containers.Sarus* attribute), 124

`workdir` (*reframe.core.containers.ContainerPlatform* attribute), 124

Z

`zip()` (in module *reframe.utility.sanity*), 131