
signac Documentation

Release 1.7.0

Carl Simon Adorf

Nov 30, 2022

Contents

1	Contents	3
1.1	Installation	3
1.2	Command Line Interface	4
1.3	API Reference	14
1.4	Changelog	87
1.5	Support and Development	108
2	Indices and tables	111
	Python Module Index	113
	Index	115

Note: This is documentation for the **signac** *core package*, which is part of the **signac** framework. See [here](#) for a comprehensive introduction to the **signac** *framework*.

1.1 Installation

The recommended installation method for **signac** is via [conda](#) or [pip](#). The software is tested for Python versions 3.6+ and has minimal dependencies. Some features such as the HDF5 integration require additional packages. Supported Python and NumPy versions are determined according to the [NEP 29 deprecation policy](#).

1.1.1 Install with conda

You can install **signac** via conda (available on the [conda-forge](#) channel), with:

```
$ conda install -c conda-forge signac
```

All additional dependencies will be installed automatically. To upgrade the package, execute:

```
$ conda update signac
```

1.1.2 Install with pip

To install the package with the package manager [pip](#), execute

```
$ pip install signac --user
```

Note: It is highly recommended to install the package into the user space and not as superuser!

To upgrade the package, simply execute the same command with the `--upgrade` option.

```
$ pip install signac --user --upgrade
```

Consider installing optional dependencies:

```
$ pip install pymongo passlib bcrypt --user
```

1.1.3 Source Code Installation

Alternatively you can clone the [git repository](#) and execute the `setup.py` script to install the package.

```
git clone https://github.com/glotzerlab/signac.git
cd signac
python setup.py install --user
```

Consider installing *optional dependencies*.

1.1.4 Optional dependencies

Unless you install via `conda`, optional dependencies are not installed automatically. In case you want to use extra features that require external packages, you need to install these manually.

Extra features with dependencies:

MongoDB database backend required: `pymongo`

recommended: `passlib, bcrypt`

HDF5 integration required: `h5py`

1.2 Command Line Interface

The following core **signac** functions are — in addition to the Python interface — accessible directly via the `$ signac` command.

The commands can be roughly grouped by task, ordered by frequency of use:

Using a project:

- *statepoint*
- *document*
- *find*
- *shell*
- *schema*
- *diff*

Modifying a project:

- *init*
- *config*
- *project*
- *job*
- *rm*
- *update-cache*

Sharing and archiving a project:

- *view*
- *export*
- *import*
- *move*
- *clone*
- *sync*
- *index*

1.2.1 clone

```
usage: signac clone [-h] project job_id [job_id ...]

positional arguments:
  project      The root directory of the project to clone one or more jobs in.
  job_id      One or more job ids. The corresponding jobs must be initialized.

optional arguments:
  -h, --help  show this help message and exit
```

1.2.2 config

```
usage: signac config [-h] [-g] [-l] [-f] {show,set,host,verify} ...

positional arguments:
  {show,set,host,verify}

optional arguments:
  -h, --help            show this help message and exit
  -g, --global          Modify the global configuration.
  -l, --local           Modify the local configuration.
  -f, --force           Skip sanity checks when modifying the configuration.
```

1.2.3 diff

```
usage: signac diff [-h] [-p [PRETTY]] [-i [INDENT]] [-f FILTER [FILTER ...]]
                  [-d DOC_FILTER [DOC_FILTER ...]]
                  [job_id [job_id ...]]

Find the difference among job state points.

positional arguments:
  job_id      One or more job ids. The corresponding jobs must be
              initialized.

optional arguments:
  -h, --help            show this help message and exit
  -p [PRETTY], --pretty [PRETTY]
```

(continues on next page)

(continued from previous page)

```

Print state point in pretty format. An optional
argument to this flag specifies the maximal depth a
state point is printed.
-i [INDENT], --indent [INDENT]
Specify the indentation of the JSON formatted state
point.
-f FILTER [FILTER ...], --filter FILTER [FILTER ...]
Limit the diff to jobs matching this state point
filter.
-d DOC_FILTER [DOC_FILTER ...], --doc-filter DOC_FILTER [DOC_FILTER ...]
Show documents of jobs matching this document filter.

```

1.2.4 document

```

usage: signac document [-h] [-p [PRETTY]] [-i [INDENT]] [-s]
[-f FILTER [FILTER ...]]
[-d DOC_FILTER [DOC_FILTER ...]] [--index INDEX]
[job_id [job_id ...]]

Print the document(s) corresponding to one or more job ids.

positional arguments:
  job_id                One or more job ids. The job corresponding to a job id
                        must be initialized.

optional arguments:
  -h, --help            show this help message and exit
  -p [PRETTY], --pretty [PRETTY]
                        Print document in pretty format. An optional argument
                        to this flag specifies the maximal depth a document is
                        printed.
  -i [INDENT], --indent [INDENT]
                        Specify the indentation of the JSON formatted state
                        point.
  -s, --sort            Sort the document keys for output in JSON format.
  -f FILTER [FILTER ...], --filter FILTER [FILTER ...]
                        Show documents of jobs matching this state point
                        filter.
  -d DOC_FILTER [DOC_FILTER ...], --doc-filter DOC_FILTER [DOC_FILTER ...]
                        Show documents of job matching this document filter.
  --index INDEX        The filename of an index file.

```

1.2.5 export

```

usage: signac export [-h] [--move] [-f FILTER [FILTER ...]]
[-d DOC_FILTER [DOC_FILTER ...]] [-j JOB_ID [JOB_ID ...]]
target [schema_path]

Export the project data space (or a subset) to a directory, a zipfile, or a
tarball.

positional arguments:

```

(continues on next page)

(continued from previous page)

target	The target to export to. May be a path to a directory, a zipfile, or a tarball.
schema_path	Specify an optional export path, based on the job state point, e.g., 'foo/{job.sp.foo}'.
optional arguments:	
-h, --help	show this help message and exit
--move	Move data to export target instead of copying. Can only be used when exporting to a directory target.
select:	
-f FILTER [FILTER ...], --filter FILTER [FILTER ...]	Limit the jobs to export to those matching the state point filter.
-d DOC_FILTER [DOC_FILTER ...], --doc-filter DOC_FILTER [DOC_FILTER ...]	Limit the jobs to export to those matching this document filter.
-j JOB_ID [JOB_ID ...], --job-id JOB_ID [JOB_ID ...]	Limit the jobs to export to those matching the provided job ids.

1.2.6 find

usage: signac find [-h] [-d DOC_FILTER [DOC_FILTER ...]] [-i INDEX] [-s [SHOW]] [--sp [SP [SP ...]]] [--doc [DOC [DOC ...]]] [-p [PRETTY]] [-1] [filter [filter ...]]	
All filter arguments may be provided either directly in JSON encoding or in a simplified form, e.g., -- \$ signac find a 42 -- is equivalent to -- \$ signac find '{"a": 42}'.	
positional arguments:	
filter	A JSON encoded state point filter (key-value pairs).
optional arguments:	
-h, --help	show this help message and exit
-d DOC_FILTER [DOC_FILTER ...], --doc-filter DOC_FILTER [DOC_FILTER ...]	A document filter.
-i INDEX, --index INDEX	The filename of an index file.
-s [SHOW], --show [SHOW]	Show the state point and document of each job. Equivalent to --sp --doc --pretty 3.
--sp [SP [SP ...]]	Show the state point of each job. Can be passed the list of state point keys to print (if they exist for a given job).
--doc [DOC [DOC ...]]	Show the document of each job. Can be passed the list of document keys to print (if they exist for a given job).
-p [PRETTY], --pretty [PRETTY]	Pretty print output when using --sp, --doc, or --show. Argument is the depth to which keys are printed.

(continues on next page)

(continued from previous page)

```
-l, --one-line      Print output in JSON and on one line.
```

1.2.7 import

```
usage: signac import [-h] [--move] [--sync] [--sync-interactive]
                   [origin] [schema_path]
```

Import an existing dataset into this project. Optionally provide a file path based schema to specify the state point metadata. Providing a path based schema is only necessary if the data set was not previously exported from a signac project.

positional arguments:

```
  origin            The origin to import from. May be a path to a directory,
                   a zipfile, or a tarball. Defaults to the current working
                   directory.
  schema_path       Specify an optional import path, such as
                   'foo/{foo:int}'. Possible type definitions include bool,
                   int, float, and str. The type is assumed to be 'str' if
                   no type is specified.
```

optional arguments:

```
-h, --help          show this help message and exit
--move             Move the data upon import instead of copying. Can only
                   be used when importing from a directory.
--sync            Attempt recursive synchronization with default
                   arguments.
--sync-interactive Synchronize the project with the origin data space
                   interactively.
```

1.2.8 index

```
usage: signac index [-h] [-t TAGS [TAGS ...]] [root]
```

positional arguments:

```
  root            Specify the root path from where the main index is to
                   be compiled.
```

optional arguments:

```
-h, --help          show this help message and exit
-t TAGS [TAGS ...], --tags TAGS [TAGS ...]
                   Specify tags for this main index compilation.
```

1.2.9 init

```
usage: signac init [-h] [-w WORKSPACE] project_id
```

positional arguments:

```
  project_id      Initialize a project with the given project id.
```

(continues on next page)

(continued from previous page)

```
optional arguments:
-h, --help          show this help message and exit
-w WORKSPACE, --workspace WORKSPACE
                    The path to the workspace directory.
```

1.2.10 job

```
usage: signac job [-h] [-w] [-c] [statepoint]

positional arguments:
statepoint          The job's statepoint in JSON format. Omit this argument to
                    read from STDIN.

optional arguments:
-h, --help          show this help message and exit
-w, --workspace    Print the job's workspace path instead of the job id.
-c, --create        Create the job's workspace directory if necessary.
```

1.2.11 move

```
usage: signac move [-h] project job_id [job_id ...]

positional arguments:
project            The root directory of the project to move one or more jobs to.
job_id            One or more job ids. The corresponding jobs must be initialized.

optional arguments:
-h, --help        show this help message and exit
```

1.2.12 project

```
usage: signac project [-h] [-w] [-i] [-a]

optional arguments:
-h, --help          show this help message and exit
-w, --workspace    Print the project's workspace path instead of the project
                    id.
-i, --index         Generate and print an index for the project.
-a, --access        Create access module for indexing.
```

1.2.13 rm

```
usage: signac rm [-h] [-c] [-i] [-v] job_id [job_id ...]

positional arguments:
job_id            One or more job ids of jobs to remove.

optional arguments:
```

(continues on next page)

(continued from previous page)

```

-h, --help          show this help message and exit
-c, --clear         Do not completely remove, but only clear the job.
-i, --interactive   Request confirmation before attempting to remove/clear
                   each job.
-v, --verbose       Be verbose when removing/clearing files.

```

1.2.14 schema

```

usage: signac schema [-h] [-x] [-t DEPTH] [-p PRECISION] [-r MAX_NUM_RANGE]
                   [-f FILTER [FILTER ...]] [-d DOC_FILTER [DOC_FILTER ...]]
                   [-j JOB_ID [JOB_ID ...]]

```

optional arguments:

```

-h, --help          show this help message and exit
-x, --exclude-const Exclude state point parameters, which are constant
                   over the complete project data space.
-t DEPTH, --depth DEPTH
                   A non-zero value will format the schema in a nested
                   representation up to the specified depth. The default
                   is a flat view (depth=0).
-p PRECISION, --precision PRECISION
                   Round all numerical values up to the given precision.
-r MAX_NUM_RANGE, --max-num-range MAX_NUM_RANGE
                   The maximum number of entries shown for a value range,
                   defaults to 5.

```

select:

```

-f FILTER [FILTER ...], --filter FILTER [FILTER ...]
                   Detect schema only for jobs that match the state point
                   filter.
-d DOC_FILTER [DOC_FILTER ...], --doc-filter DOC_FILTER [DOC_FILTER ...]
                   Detect schema only for jobs that match the document
                   filter.
-j JOB_ID [JOB_ID ...], --job-id JOB_ID [JOB_ID ...]
                   Detect schema only for jobs with the given job ids.

```

1.2.15 shell

```

usage: signac shell [-h] [-c COMMAND] [-f FILTER [FILTER ...]]
                   [-d DOC_FILTER [DOC_FILTER ...]] [-j JOB_ID [JOB_ID ...]]
                   [file]

```

positional arguments:

```

file              Execute Python script in file.

```

optional arguments:

```

-h, --help          show this help message and exit
-c COMMAND, --command COMMAND
                   Execute Python program passed as string.

```

select:

```

Specify one or more jobs to preset the `jobs` variable as a generator over

```

(continues on next page)

(continued from previous page)

```

all job handles associated with the given selection. If the selection
contains only one job, an additional `job` variable is referencing that
single job, otherwise it is `None`.

-f FILTER [FILTER ...], --filter FILTER [FILTER ...]
    Reduce selection to jobs that match the given filter.
-d DOC_FILTER [DOC_FILTER ...], --doc-filter DOC_FILTER [DOC_FILTER ...]
    Reduce selection to jobs that match the given document
    filter.
-j JOB_ID [JOB_ID ...], --job-id JOB_ID [JOB_ID ...]
    Reduce selection to jobs that match the given job ids.

```

1.2.16 statepoint

```

usage: signac statepoint [-h] [-p [PRETTY]] [-i [INDENT]] [-s]
                        [job_id [job_id ...]]

Print the statepoint(s) corresponding to one or more job ids.

positional arguments:
  job_id                One or more job ids. The corresponding jobs must be
                        initialized.

optional arguments:
  -h, --help            show this help message and exit
  -p [PRETTY], --pretty [PRETTY]
                        Print state point in pretty format. An optional
                        argument to this flag specifies the maximal depth a
                        state point is printed.
  -i [INDENT], --indent [INDENT]
                        Specify the indentation of the JSON formatted state
                        point.
  -s, --sort            Sort the state point keys for output.

```

1.2.17 sync

```

usage: signac sync [-h] [-v] [-a] [-r] [-l] [-p] [-o] [-g] [-t] [-x [EXCLUDE]]
                  [-I] [--size-only] [--round-times] [-n] [-u]
                  [-s {always,never,update,Ask}] [-k KEY] [--all-keys]
                  [--no-keys] [-w] [--force] [-m] [--parallel [PARALLEL]]
                  [--stats] [-H] [--json] [-f FILTER [FILTER ...]]
                  [-d DOC_FILTER [DOC_FILTER ...]] [-j JOB_ID [JOB_ID ...]]
                  source [destination]

```

Use this command to synchronize this project with another project; similar to the synchronization of two directories with `rsync`. Data is always copied from the source to the destination. For example: `signac sync /path/to/other/project -u --all-keys` means "Synchronize all jobs within this project with those in the other project; overwrite files if the source files is newer and overwrite all conflicting keys in the project and job documents."

positional arguments:

(continues on next page)

(continued from previous page)

```

source          The root directory of the project that this project
                should be synchronized with.
destination     Optional: The root directory of the project that
                should be modified for synchronization, defaults to
                the local project.

optional arguments:
-h, --help      show this help message and exit
-v, --verbosity Set level of verbosity.
-w, --allow-workspace
                Allow the specification of a workspace (instead of a
                project) directory as the destination path.
--force         Ignore all warnings, just synchronize.
-m, --merge     Clone all the jobs that are not present in destination
                from source.
--parallel [PARALLEL]
                Use multiple threads for synchronization. You may
                optionally specify how many threads to use, otherwise
                all available processing units will be utilized.
--stats         Provide file transfer statistics.
-H, --human-readable
                Provide statistics with human readable formatting.
--json         Print statistics in JSON formatting.

copy options:
-a, --archive   archive mode; equivalent to: '-rltpog'
-r, --recursive Do not skip sub-directories, but synchronize
                recursively.
-l, --links     Copy symbolic links as symbolic links pointing to the
                original source.
-p, --perms    Preserve permissions.
-o, --owner    Preserve owner.
-g, --group    Preserve group.
-t, --times    Preserve file modification times (requires -p).
-x [EXCLUDE], --exclude [EXCLUDE]
                Exclude all files matching the given pattern. Exclude
                all files if this option is provided without any
                argument.
-I, --ignore-times
                Never rely on file meta data such as the size or the
                modification time when determining file differences.
--size-only     Ignore modification times during file comparison.
                Useful when synchronizing between file systems with
                different timestamp resolution.
--round-times   Round modification times during file comparison.
                Useful when synchronizing between file systems with
                different timestamp resolution.
-n, --dry-run   Do not actually execute the synchronization. Increase
                the output verbosity to see messages about what would
                potentially happen.
-u, --update    Skip files with newer modification time stamp. This is
                a short-cut for: --strategy=update.

sync strategy:
-s {always,never,update,Ask}, --strategy {always,never,update,Ask}
                Specify a synchronization strategy, for differing
                files.
-k KEY, --key KEY
                Specify a regular expression for keys that should be
                overwritten as part of the project and job document

```

(continues on next page)

(continued from previous page)

```

synchronization.
--all-keys      Overwrite all conflicting keys. Equivalent to
                '--key=*. *'.
--no-keys       Never overwrite any conflicting keys.

select:
-f FILTER [FILTER ...], --filter FILTER [FILTER ...]
                Only synchronize jobs that match the state point
                filter.
-d DOC_FILTER [DOC_FILTER ...], --doc-filter DOC_FILTER [DOC_FILTER ...]
                Only synchronize jobs that match the document filter.
-j JOB_ID [JOB_ID ...], --job-id JOB_ID [JOB_ID ...]
                Only synchronize jobs with the given job ids.

```

1.2.18 update-cache

```

usage: signac update-cache [-h]

Use this command to update the project's persistent state point cache. This
feature is still experimental and may be removed in future versions.

optional arguments:
-h, --help  show this help message and exit

```

1.2.19 view

```

usage: signac view [-h] [-f FILTER [FILTER ...]]
                 [-d DOC_FILTER [DOC_FILTER ...]] [-j JOB_ID [JOB_ID ...]]
                 [-i INDEX]
                 [prefix] [path]

Generate a human readable set of paths representing state points in the
workspace, e.g.
view/param_name_1/param_value_1/param_name_2/param_value_2/job. The leaf nodes
of this directory structure are symlinks (named "job") into the workspace
directory for that parameter combination. Note that all positional arguments
must be provided before any keyword arguments. In particular, the prefix and
path must be specified before arguments such as the filters, e.g. signac view
$PREFIX $VIEW_PATH -f FILTERS -d DOC_FILTERS.

positional arguments:
prefix                The path where the view is to be created. Defaults to
                    view.
path                 The path used for the generation of the linked view
                    hierarchy, defaults to '{{auto}}' (see
                    Project.export_to for information on how this is
                    expanded).

optional arguments:
-h, --help            show this help message and exit

select:

```

(continues on next page)

(continued from previous page)

```
-f FILTER [FILTER ...], --filter FILTER [FILTER ...]
    Limit the view to jobs matching this state point
    filter.
-d DOC_FILTER [DOC_FILTER ...], --doc-filter DOC_FILTER [DOC_FILTER ...]
    Limit the view to jobs matching this document filter.
-j JOB_ID [JOB_ID ...], --job-id JOB_ID [JOB_ID ...]
    Limit the view to jobs with these job ids.
-i INDEX, --index INDEX
    The filename of an index file.
```

1.3 API Reference

This is the API for the **signac** (core) application.

1.3.1 The Project

Attributes

<code>Project.build_job_search_index(index[, _trust])</code>	Build a job search index.
<code>Project.build_job_statepoint_index([...])</code>	Build a state point index to identify jobs with specific parameters.
<code>Project.check()</code>	Check the project's workspace for corruption.
<code>Project.clone(job[, copytree])</code>	Clone job into this project.
<code>Project.config</code>	Get project's configuration.
<code>Project.create_access_module([filename, ...])</code>	Create the access module for indexing.
<code>Project.create_linked_view([prefix, ...])</code>	Create or update a persistent linked view of the selected data space.
<code>Project.detect_schema([exclude_const, ...])</code>	Detect the project's state point schema.
<code>Project.data</code>	Get data associated with this project.
<code>Project.doc</code>	Get document associated with this project.
<code>Project.document</code>	Get document associated with this project.
<code>Project.dump_statepoints(statepoints)</code>	Dump the state points and associated job ids.
<code>Project.export_to(target[, path, copytree])</code>	Export all jobs to a target location, such as a directory or a (compressed) archive file.
<code>Project.find_job_ids([filter, doc_filter, index])</code>	Find the job_ids of all jobs matching the filters.
<code>Project.find_jobs([filter, doc_filter])</code>	Find all jobs in the project's workspace.
<code>Project.fn(filename)</code>	Prepend a filename with the project's root directory path.
<code>Project.get_id()</code>	Get the project identifier.
<code>Project.get_statepoint(jobid[, fn])</code>	Get the state point associated with a job id.
<code>Project.groupby([key, default])</code>	Group jobs according to one or more state point parameters.
<code>Project.groupbydoc([key, default])</code>	Group jobs according to one or more document values.
<code>Project.import_from([origin, schema, sync, ...])</code>	Import the data space located at origin into this project.
<code>Project.id</code>	Get the project identifier.

Continued on next page

Table 1 – continued from previous page

<code>Project.index([formats, depth, skip_errors, ...])</code>	Generate an index of the project's workspace.
<code>Project.isfile(filename)</code>	Check if a filename exists in the project's root directory.
<code>Project.min_len_unique_id()</code>	Determine the minimum length required for a job id to be unique.
<code>Project.num_jobs()</code>	Return the number of initialized jobs.
<code>Project.open_job([statepoint, id])</code>	Get a job handle associated with a state point.
<code>Project.read_statepoints([fn])</code>	Read all state points from a file.
<code>Project.repair([fn_statepoints, index, job_ids])</code>	Attempt to repair the workspace after it got corrupted.
<code>Project.reset_statepoint(job, new_statepoint)</code>	Overwrite the state point of this job while preserving job data.
<code>Project.root_directory()</code>	Return the project's root directory.
<code>Project.stores</code>	Get HDF5-stores associated with this project.
<code>Project.sync(other[, strategy, exclude, ...])</code>	Synchronize this project with the other project.
<code>Project.update_cache()</code>	Update the persistent state point cache.
<code>Project.update_statepoint(job, update[, ...])</code>	Change the state point of this job while preserving job data.
<code>Project.workspace()</code>	Return the project's workspace directory.
<code>Project.write_statepoints([statepoints, fn, ...])</code>	Dump state points to a file.

class `signac.Project` (*config=None, ignore_schema_version=False*)

Bases: `object`

The handle on a signac project.

Application developers should usually not need to directly instantiate this class, but use `get_project()` instead.

Parameters

- **config** – The project configuration to use. By default, it loads the first signac project configuration found while searching upward from the current working directory (Default value = None).
- **ignore_schema_version** (*bool*) – (Default value = False).

FN_CACHE = `'signac_sp_cache.json.gz'`

The default filename for the state point cache file.

FN_DOCUMENT = `'signac_project_document.json'`

The project's document filename.

FN_STATEPOINTS = `'signac_statepoints.json'`

The default filename to read from and write state points to.

KEY_DATA = `'signac_data'`

The project's datastore key.

build_job_search_index (*index, trust=False*)

Build a job search index.

Parameters

- **index** (*list*) – A document index.
- **trust** – (Default value = False).

Returns A job search index based on the provided index.

Return type `JobSearchIndex`

Deprecated since version 1.3: This will be removed in 2.0.

build_job_statepoint_index (*exclude_const=False, index=None*)

Build a state point index to identify jobs with specific parameters.

This method generates pairs of state point keys and mappings of values to a set of all corresponding job ids. The pairs are ordered by the number of different values. Since state point keys may be nested, they are represented as a tuple. For example:

```
>>> for i in range(4):
...     project.open_job({'a': i, 'b': {'c': i % 2}}).init()
...
>>> for key, value in project.build_job_statepoint_index():
...     print(key)
...     pprint.pprint(value)
...
('b', 'c')
defaultdict(<class 'set'>,
            {0: {'3a530c13bfaf57517b4e81ecab6aec7f',
                 '4e9a45a922eae6bb5d144b36d82526e4'},
             1: {'d49c6609da84251ab096654971115d0c',
                 '5c2658722218d48a5eb1e0ef7c26240b'}})
('a',)
defaultdict(<class 'set'>,
            {0: {'4e9a45a922eae6bb5d144b36d82526e4'},
             1: {'d49c6609da84251ab096654971115d0c'},
             2: {'3a530c13bfaf57517b4e81ecab6aec7f'},
             3: {'5c2658722218d48a5eb1e0ef7c26240b'}})
```

Values that are constant over the complete data space can be optionally ignored with the *exclude_const* argument set to True.

Parameters

- **exclude_const** (*bool*) – Exclude entries that are shared by all jobs that are part of the index (Default value = False).
- **index** – A document index.

Yields *tuple* – Pairs of state point keys and mappings of values to a set of all corresponding job ids (Default value = None).

Deprecated since version 1.3: This will be removed in 2.0. Use the `detect_schema()` function instead.

check ()

Check the project's workspace for corruption.

Raises `signac.errors.JobsCorruptedError` – When one or more jobs are identified as corrupted.

clone (*job, copytree=<function copytree>*)

Clone job into this project.

Create an identical copy of job within this project.

See `signac clone` for the command line equivalent.

Parameters

- **job** (*Job*) – The job to copy into this project.
- **copytree** – (Default value = `shutil.copytree()`)

Returns The job instance corresponding to the copied job.

Return type *Job*

Raises *DestinationExistsError* – In case that a job with the same id is already initialized within this project.

config

Get project's configuration.

Returns Dictionary containing project's configuration.

Return type *_ProjectConfig*

create_access_module (*filename=None, main=True, master=None*)

Create the access module for indexing.

This method generates the access module required to make this project's index part of a main index.

Parameters

- **filename** (*str*) – The name of the access module file. Defaults to the standard name and should usually not be changed.
- **main** (*bool*) – If True, add directives for the compilation of a master index when executing the module (Default value = True).
- **master** (*bool*) – Deprecated parameter. Replaced by main.

Returns Access module name.

Return type *str*

Deprecated since version 1.5: This will be removed in 2.0. Access modules are deprecated.

create_linked_view (*prefix=None, job_ids=None, index=None, path=None*)

Create or update a persistent linked view of the selected data space.

Similar to *export_to()*, this function expands the data space for the selected jobs, but instead of copying data will create symbolic links to the individual job workspace directories. This is primarily useful for browsing through the data space using a file-browser with human-interpretable directory paths.

By default, the paths of the view will be based on variable state point keys as part of the *implicit* schema of the selected jobs that we create the view for. For example, creating a linked view for a data space with schema

```
>>> print(project.detect_schema())
{
  'foo': 'int([0, 1, 2, ..., 8, 9], 10)',
}
```

by calling `project.create_linked_view('my_view')` will look similar to:

```
my_view/foo/0/job -> workspace/b8fcc6b8f99c56509eb65568922e88b8
my_view/foo/1/job -> workspace/b6cd26b873ae3624653c9268deff4485
...
```

It is possible to control the paths using the `path` argument, which behaves in the exact same manner as the equivalent argument for *export_to()*.

Note: The behavior of this function is almost equivalent to `project.export_to('my_view', copytree=os.symlink)` with the major difference that view hierarchies are actually *updated*, meaning that invalid links are automatically removed.

See *signac view* for the command line equivalent.

Parameters

- **prefix** (*str*) – The path where the linked view will be created or updated (Default value = None).
- **job_ids** (*iterable*) – If None (the default), create the view for the complete data space, otherwise only for this iterable of job ids.
- **index** – A document index (Default value = None).
- **path** – The path (function) used to structure the linked data space (Default value = None).

Returns A dictionary that maps the source directory paths to the linked directory paths.

Return type *dict*

data

Get data associated with this project.

This property should be used for large array-like data, which can't be stored efficiently in the project document. For examples and usage, see [Centralized Project Data](#).

Equivalent to:

```
return project.stores['signac_data']
```

See also:

H5Store: Usage examples.

Returns An HDF5-backed datastore.

Return type *H5Store*

detect_schema (*exclude_const=False, subset=None, index=None*)

Detect the project's state point schema.

See *signac schema* for the command line equivalent.

Parameters

- **exclude_const** (*bool*) – Exclude all state point keys that are shared by all jobs within this project (Default value = False).
- **subset** – A sequence of jobs or job ids specifying a subset over which the state point schema should be detected (Default value = None).
- **index** – A document index (Default value = None).

Returns The detected project schema.

Return type *ProjectSchema*

doc

Get document associated with this project.

Alias for *document* ().

Returns The project document.

Return type *BufferedJSONAttrDict*

document

Get document associated with this project.

Returns The project document.

Return type *BufferedJSONAttrDict*

dump_statepoints (*statepoints*)

Dump the state points and associated job ids.

Equivalent to:

```
{project.open_job(sp).id: sp for sp in statepoints}
```

Parameters **statepoints** (*iterable*) – A list of state points.

Returns A mapping, where the key is the job id and the value is the state point.

Return type *dict*

export_to (*target, path=None, copytree=None*)

Export all jobs to a target location, such as a directory or a (compressed) archive file.

Use this function in combination with *find_jobs()* to export only a select number of jobs, for example:

```
project.find_jobs({'foo': 0}).export_to('foo_0.tar')
```

The *path* argument enables users to control how exactly the exported data space is to be expanded. By default, the path-function will be based on the *implicit* schema of the exported jobs. For example, exporting jobs that all differ by a state point key *foo* with `project.export_to('data/')`, the exported directory structure could look like this:

```
data/foo/0
data/foo/1
...
```

That would be equivalent to specifying `path=lambda job: os.path.join('foo', job.sp.foo)`.

Instead of a function, we can also provide a string, where fields for state point keys are automatically formatted. For example, the following two path arguments are equivalent: “foo/{foo}” and “foo/{job.sp.foo}”.

Any attribute of job can be used as a field here, so `job.doc.bar`, `job.id`, and `job.ws` can also be used as path fields.

A special `{{auto}}` field allows us to expand the path automatically with state point keys that have not been specified explicitly. So, for example, one can provide `path="foo/{foo}/{{auto}}"` to specify that the path shall begin with `foo/{foo}/`, but is then automatically expanded with all other state point key-value pairs. How key-value pairs are concatenated can be controlled *via* the format-specifier, so for example, `path="{{auto: _}}"` will generate a structure such as

```
data/foo_0
data/foo_1
...
```

Finally, providing `path=False` is equivalent to `path="{job.id}"`.

See also:

import_from() : Previously exported or non-signac data spaces can be imported.

signac export : See *signac export* for the command line equivalent.

Parameters

- **target** – A path to a directory to export to. The target can not already exist. Besides directories, possible targets are tar files (*.tar*), gzipped tar files (*.tar.gz*), zip files (*.zip*), bzip2-compressed files (*.bz2*), and xz-compressed files (*.xz*).
- **path** – The path (function) used to structure the exported data space. This argument must either be a callable which returns a path (str) as a function of *job*, a string where fields are replaced using the job-state point dictionary, or *False*, which means that we just use the job-id as path. Defaults to the equivalent of `{{auto}}`.
- **copytree** – The function used for the actual copying of directory tree structures. Defaults to `shutil.copytree()`. Can only be used when the target is a directory.

Returns A dict that maps the source directory paths, to the target directory paths.

Return type dict

find_job_ids (*filter=None, doc_filter=None, index=None*)

Find the job_ids of all jobs matching the filters.

The optional filter arguments must be a Mapping of key-value pairs and JSON serializable.

Note: Providing a pre-calculated index may vastly increase the performance of this function.

Parameters

- **filter** (*Mapping*) – A mapping of key-value pairs that all indexed job state points are compared against (Default value = None).
- **doc_filter** (*Mapping*) – A mapping of key-value pairs that all indexed job documents are compared against (Default value = None).
- **index** – A document index. If not provided, an index will be computed (Default value = None).

Returns

Return type The ids of all indexed jobs matching both filters.

Raises

- `TypeError` – If the filters are not JSON serializable.
- `ValueError` – If the filters are invalid.
- `RuntimeError` – If the filters are not supported by the index.

Deprecated since version 1.3: This will be removed in 2.0. Use `find_jobs()` instead, then access ids with `job.id`. Replicate the original behavior with `[job.id for job in project.find_jobs()]`

find_jobs (*filter=None, doc_filter=None*)

Find all jobs in the project's workspace.

The optional filter arguments must be a Mapping of key-value pairs and JSON serializable. The *filter* argument is used to search against job state points, whereas the *doc_filter* argument compares against job document keys.

See [signac find](#) for the command line equivalent.

Parameters

- **filter** (*Mapping*) – A mapping of key-value pairs that all indexed job state points are compared against (Default value = None).
- **doc_filter** (*Mapping*) – A mapping of key-value pairs that all indexed job documents are compared against (Default value = None).

Returns JobsCursor of jobs matching the provided filter(s).

Return type *JobsCursor*

Raises

- `TypeError` – If the filters are not JSON serializable.
- `ValueError` – If the filters are invalid.
- `RuntimeError` – If the filters are not supported by the index.

fn (*filename*)

Prepend a filename with the project's root directory path.

Parameters **filename** (*str*) – The name of the file.

Returns The joined path of project root directory and filename.

Return type *str*

get_id ()

Get the project identifier.

Returns The project id.

Return type *str*

Deprecated since version 1.3: This will be removed in 2.0. Use project.id instead.

classmethod **get_job** (*root=None*)

Find a Job in or above the current working directory (or provided path).

Parameters **root** (*str*) – The job root directory. If no root directory is given, the current working directory is assumed to be the job directory (Default value = None).

Returns The job instance.

Return type *Job*

Raises `LookupError` – When job cannot be found.

classmethod **get_project** (*root=None, search=True, **kwargs*)

Find a project configuration and return the associated project.

Parameters

- **root** (*str*) – The starting point to search for a project, defaults to the current working directory.
- **search** (*bool*) – If True, search for project configurations inside and above the specified root directory, otherwise only return projects with a root directory identical to the specified root argument (Default value = True).
- ****kwargs** – Optional keyword arguments that are forwarded to the *Project* class constructor.

Returns An instance of *Project*.

Return type *Project*

Raises `LookupError` – When project configuration cannot be found.

get_statepoint (*jobid*, *fn=None*)

Get the state point associated with a job id.

The state point is retrieved from the internal cache, from the workspace or from a state points file.

Parameters

- **jobid** (*str*) – A job id to get the state point for.
- **fn** (*str*) – The filename of the file containing the state points, defaults to `FN_STATEPOINTS`.

Returns The state point corresponding to jobid.

Return type `dict`

Raises

- `KeyError` – If the state point associated with jobid could not be found.
- `signac.errors.JobsCorruptedError` – If the state point manifest file corresponding to jobid is inaccessible or corrupted.

Deprecated since version 1.3: This will be removed in 2.0. Use `open_job(id=jobid).statepoint()` function instead.

groupby (*key=None*, *default=None*)

Group jobs according to one or more state point parameters.

This method can be called on any `JobsCursor` such as the one returned by `find_jobs()` or by iterating over a project.

Examples

```
# Group jobs by state point parameter 'a'.
for key, group in project.groupby('a'):
    print(key, list(group))

# Group jobs by document value 'a'.
for key, group in project.groupby('doc.a'):
    print(key, list(group))

# Group jobs by jobs.sp['a'] and job.document['b']
for key, group in project.groupby('a', 'doc.b'):
    print(key, list(group))

# Find jobs where job.sp['a'] is 1 and group them
# by job.sp['b'] and job.sp['c'].
for key, group in project.find_jobs({'a': 1}).groupby(('b', 'c')):
    print(key, list(group))

# Group by job.sp['d'] and job.document['count'] using a lambda.
for key, group in project.groupby(
    lambda job: (job.sp['d'], job.document['count'])
):
    print(key, list(group))
```

If *key* is `None`, jobs are grouped by id, placing one job into each group.

Parameters

- **key** (*str*, *iterable*, or *callable*) – The state point grouping parameter(s) passed as a string, iterable of strings, or a callable that will be passed one argument, the job (Default value = None).
- **default** – A default value to be used when a given state point key is not present. The value must be sortable and is only used if not None (Default value = None).

Returns

- **key** (*str*) – Grouped key.
- **group** (*iterable of Jobs*) – Iterable of *Job* instances matching this group key.

groupbydoc (*key=None*, *default=None*)

Group jobs according to one or more document values.

This method can be called on any *JobsCursor* such as the one returned by *find_jobs()* or by iterating over a project.

Examples

```
# Group jobs by document value 'a'.
for key, group in project.groupbydoc('a'):
    print(key, list(group))

# Find jobs where job.sp['a'] is 1 and group them
# by job.document['b'] and job.document['c'].
for key, group in project.find_jobs({'a': 1}).groupbydoc(('b', 'c')):
    print(key, list(group))

# Group by whether 'd' is a field in the job.document using a lambda.
for key, group in project.groupbydoc(lambda doc: 'd' in doc):
    print(key, list(group))
```

If *key* is None, jobs are grouped by id, placing one job into each group.

Parameters

- **key** (*str*, *iterable*, or *callable*) – The document grouping parameter(s) passed as a string, iterable of strings, or a callable that will be passed one argument, *document* (Default value = None).
- **default** – A default value to be used when a given document key is not present. The value must be sortable and is only used if not None (Default value = None).

Deprecated since version 1.7: This will be removed in 2.0. Use *groupby* with a ‘doc.’ filter instead, see <https://docs.signac.io/en/latest/query.html#query-namespaces>.

id

Get the project identifier.

Returns The project id.

Return type *str*

import_from (*origin=None*, *schema=None*, *sync=None*, *copytree=None*)

Import the data space located at *origin* into this project.

This function will walk through the data space located at *origin* and will try to identify data space paths that can be imported as a job workspace into this project.

The `schema` argument expects a function that takes a path argument and returns a state point dictionary. A default function is used when no argument is provided. The default schema function will simply look for state point manifest files—usually named `signac_statepoint.json`—and then import all data located within that path into the job workspace corresponding to the state point specified in the manifest file.

Alternatively the `schema` argument may be a string, that is converted into a schema function, for example: Providing `foo/{foo:int}` as `schema` argument means that all directories under `foo/` will be imported and their names will be interpreted as the value for `foo` within the state point.

Tip: Use `copytree=os.replace` or `copytree=shutil.move` to move dataspace on import instead of copying them.

Warning: Imports can fail due to conflicts. Moving data instead of copying may therefore lead to inconsistent states and users are advised to apply caution.

See also:

`export_to()` : Export the project data space.

signac import : See `signac import` for the command line equivalent.

Parameters

- **origin** – The path to the data space origin, which is to be imported. This may be a path to a directory, a zip file, or a tarball archive (Default value = None).
- **schema** – An optional schema function, which is either a string or a function that accepts a path as its first and only argument and returns the corresponding state point as dict. (Default value = None).
- **sync** – If `True`, the project will be synchronized with the imported data space. If a dict of keyword arguments is provided, the arguments will be used for `sync()` (Default value = None).
- **copytree** – Specify which exact function to use for the actual copytree operation. Defaults to `shutil.copytree()`.

Returns A dict that maps the source directory paths to the target directory paths.

Return type dict

index (*formats=None, depth=0, skip_errors=False, include_job_document=True*)

Generate an index of the project's workspace.

This generator function indexes every file in the project's workspace until the specified *depth*. The job document if it exists, is always indexed, other files need to be specified with the *formats* argument.

See `signac project -i` for the command line equivalent.

```
for doc in project.index({r'.*\.txt', 'TextFile'}):
    print(doc)
```

Parameters

- **formats** (*str, dict*) – The format definitions as a pattern string (e.g. `r'.*\.txt'`) or a mapping from pattern strings to formats (e.g. `'TextFile'`). If `None`, only the job document is indexed (Default value = None).
- **depth** (*int*) – Specifies the crawling depth. A value of 0 means no limit (Default value = 0).

- **skip_errors** (*bool*) – Skip all errors which occur during indexing. This is useful when trying to repair a broken workspace (Default value = False).
- **include_job_document** (*bool*) – Include the contents of job documents (Default value = True).

Yields *dict* – Index document.

classmethod **init_project** (*name*, *root=None*, *workspace=None*, *make_dir=True*)

Initialize a project with the given name.

It is safe to call this function multiple times with the same arguments. However, a *RuntimeError* is raised if an existing project configuration would conflict with the provided initialization parameters.

See *signac init* for the command line equivalent.

Parameters

- **name** (*str*) – The name of the project to initialize.
- **root** (*str*) – The root directory for the project. Defaults to the current working directory.
- **workspace** (*str*) – The workspace directory for the project. Defaults to a subdirectory *workspace* in the project root.
- **make_dir** (*bool*) – Create the project root directory if it does not exist yet (Default value = True).

Returns Initialized project, an instance of *Project*.

Return type *Project*

Raises *RuntimeError* – If the project root path already contains a conflicting project configuration.

isfile (*filename*)

Check if a filename exists in the project’s root directory.

Parameters **filename** (*str*) – The name of the file.

Returns True if filename exists in the project’s root directory.

Return type *bool*

min_len_unique_id ()

Determine the minimum length required for a job id to be unique.

This method’s runtime scales with the number of jobs in the workspace.

Returns Minimum string length of a unique job identifier.

Return type *int*

num_jobs ()

Return the number of initialized jobs.

Returns Count of initialized jobs.

Return type *int*

open_job (*statepoint=None*, *id=None*)

Get a job handle associated with a state point.

This method returns the job instance associated with the given state point or job id. Opening a job by a valid state point never fails. Opening a job by id requires a lookup of the state point from the job id, which may fail if the job was not previously initialized.

Parameters

- **statepoint** (*dict*) – The job’s unique set of state point parameters (Default value = None).
- **id** (*str*) – The job id (Default value = None).

Returns The job instance.

Return type *Job*

Raises

- `KeyError` – If the attempt to open the job by id fails.
- `LookupError` – If the attempt to open the job by an abbreviated id returns more than one match.

read_statepoints (*fn=None*)

Read all state points from a file.

See also:

dump_statepoints() Dump the state points and associated job ids.

write_statepoints() Dump state points to a file.

Parameters **fn** (*str*) – The filename of the file containing the state points, defaults to *FN_STATEPOINTS*.

Returns State points.

Return type *dict*

repair (*fn_statepoints=None, index=None, job_ids=None*)

Attempt to repair the workspace after it got corrupted.

This method will attempt to repair lost or corrupted job state point manifest files using a state points file or a document index or both.

Parameters

- **fn_statepoints** (*str*) – The filename of the file containing the state points, defaults to *FN_STATEPOINTS*.
- **index** – A document index (Default value = None).
- **job_ids** – An iterable of job ids that should get repaired. Defaults to all jobs.

Raises *signac.errors.JobsCorruptedError* – When one or more corrupted job could not be repaired.

reset_statepoint (*job, new_statepoint*)

Overwrite the state point of this job while preserving job data.

This method will change the job id if the state point has been altered.

Danger: Use this function with caution! Resetting a job’s state point, may sometimes be necessary, but can possibly lead to incoherent data spaces.

Parameters

- **job** (*Job*) – The job that should be reset to a new state point.

- **new_statepoint** (*mapping*) – The job’s new state point.

Raises

- *DestinationExistsError* – If a job associated with the new state point is already initialized.
- *OSError* – If the move failed due to an unknown system related error.

Deprecated since version 1.3: This will be removed in 2.0. Use `job.reset_statepoint()` instead.

root_directory ()

Return the project’s root directory.

Returns Path of project directory.

Return type `str`

stores

Get HDF5-stores associated with this project.

Use this property to access an HDF5 file within the project’s root directory using the `H5Store` dict-like interface.

This is an example for accessing an HDF5 file called `'my_data.h5'` within the project’s root directory:

```
project.stores['my_data']['array'] = np.random((32, 4))
```

This is equivalent to:

```
H5Store(project.fn('my_data.h5'))['array'] = np.random((32, 4))
```

Both the `project.stores` and the `H5Store` itself support attribute access. The above example could therefore also be expressed as:

```
project.stores.my_data.array = np.random((32, 4))
```

Returns The HDF5-Store manager for this project.

Return type `H5StoreManager`

sync (*other*, *strategy=None*, *exclude=None*, *doc_sync=None*, *selection=None*, ***kwargs*)

Synchronize this project with the other project.

Try to clone all jobs from the other project to this project. If a job is already part of this project, try to synchronize the job using the optionally specified strategies.

See `signac sync` for the command line equivalent.

Parameters

- **other** (*Project*) – The other project to synchronize this project with.
- **strategy** – A file synchronization strategy (Default value = `None`).
- **exclude** – Files with names matching the given pattern will be excluded from the synchronization (Default value = `None`).
- **doc_sync** – The function applied for synchronizing documents (Default value = `None`).
- **selection** – Only sync the given jobs (Default value = `None`).
- ****kwargs** – This method also accepts the same keyword arguments as the `sync_projects()` function.

Raises

- *DocumentSyncConflict* – If there are conflicting keys within the project or job documents that cannot be resolved with the given strategy or if there is no strategy provided.
- *FileSyncConflict* – If there are differing files that cannot be resolved with the given strategy or if no strategy is provided.
- *SchemaSyncConflict* – In case that the `check_schema` argument is `True` and the detected state point schema of this and the other project differ.

temporary_project (*name=None, dir=None*)

Context manager for the initialization of a temporary project.

The temporary project is by default created within the root project's workspace to ensure that they share the same file system. This is an example for how this method can be used for the import and synchronization of external data spaces.

```
with project.temporary_project() as tmp_project:
    tmp_project.import_from('/data')
    project.sync(tmp_project)
```

Parameters

- **name** (*str*) – An optional name for the temporary project. Defaults to a unique random string.
- **dir** (*str*) – Optionally specify where the temporary project root directory is to be created. Defaults to the project's workspace directory.

Returns An instance of *Project*.**Return type** *Project***to_dataframe** (**args, **kwargs*)Export the project metadata to a pandas *DataFrame*.The arguments to this function are forwarded to *to_dataframe()*.**Parameters**

- ***args** – Forwarded to *to_dataframe()*.
- ****kwargs** – Forwarded to *to_dataframe()*.

Returns**Return type** *DataFrame***update_cache** ()

Update the persistent state point cache.

This function updates a persistent state point cache, which is stored in the project root directory. Most data space operations, including iteration and filtering or selection are expected to be significantly faster after calling this function, especially for large data spaces.

update_statepoint (*job, update, overwrite=False*)

Change the state point of this job while preserving job data.

By default, this method will not change existing parameters of the state point of the job.

This method will change the job id if the state point has been altered.

Warning: While appending to a job's state point is generally safe, modifying existing parameters may lead to data inconsistency. Use the `overwrite` argument with caution!

Parameters

- **job** (*Job*) – The job whose state point shall be updated.
- **update** (*mapping*) – A mapping used for the state point update.
- **overwrite** (*bool, optional*) – If False, an error will be raised if the update modifies the values of existing keys in the state point. If True, any existing keys will be overwritten in the same way as `dict.update()`. Use with caution! (Default value = False).

Raises

- `KeyError` – If the update contains keys which are already part of the job's state point and `overwrite` is False.
- `DestinationExistsError` – If a job associated with the new state point is already initialized.
- `OSError` – If the move failed due to an unknown system related error.

Deprecated since version 1.3: This will be removed in 2.0. Use `job.update_statepoint()` instead.

`workspace()`

Return the project's workspace directory.

The workspace defaults to `project_root/workspace`. Configure this directory with the 'workspace_dir' attribute. If the specified directory is a relative path, the absolute path is relative from the project's root directory.

Note: The configuration will respect environment variables, such as `$HOME`.

See `signac project -w` for the command line equivalent.

Returns Path of workspace directory.

Return type `str`

`write_statepoints(statepoints=None, fn=None, indent=2)`

Dump state points to a file.

If the file already contains state points, all new state points will be appended, while the old ones are preserved.

See also:

`dump_statepoints()` Dump the state points and associated job ids.

Parameters

- **statepoints** (*iterable*) – A list of state points, defaults to all state points which are defined in the workspace.
- **fn** (*str*) – The filename of the file containing the state points, defaults to `FN_STATEPOINTS`.
- **indent** (*int*) – Specify the indentation of the JSON file (Default value = 2).

1.3.2 The JobsCursor class

Attributes

<code>JobsCursor.export_to(target[, path, copytree])</code>	Export all jobs to a target location, such as a directory or a (zipped) archive file.
<code>JobsCursor.groupby([key, default])</code>	Group jobs according to one or more state point parameters.
<code>JobsCursor.groupbydoc([key, default])</code>	Group jobs according to one or more document values.
<code>JobsCursor.to_dataframe([sp_prefix, ...])</code>	Convert the selection of jobs to a pandas <code>DataFrame</code> .

class `signac.contrib.project.JobsCursor` (*project, filter=None, doc_filter=None*)

Bases: `object`

An iterator over a search query result.

Application developers should not directly instantiate this class, but use `find_jobs()` instead.

Enables simple iteration and grouping operations.

Parameters

- **project** (*Project*) – Project handle.
- **filter** (*Mapping*) – A mapping of key-value pairs that all indexed job state points are compared against (Default value = None).

Notes

Iteration is performed by acquiring job ids from the project using `Project._find_job_ids()`. When no filter (`filter = None`) is provided, that method can take a much faster execution path, so not passing a filter (or passing `None` explicitly) to this constructor is strongly recommended over passing an empty filter (`filter = {}`) when iterating over the entire data space.

export_to (*target, path=None, copytree=None*)

Export all jobs to a target location, such as a directory or a (zipped) archive file.

See also:

`export_to()` : For full details on how to use this function.

Parameters

- **target** (*str*) – A path to a directory or archive file to export to.
- **path** (*str or callable*) – The path (function) used to structure the exported data space (Default value = None).
- **copytree** (*callable*) – The function used for copying of directory tree structures. Defaults to `shutil.copytree()`. Can only be used when the target is a directory (Default value = None).

Returns A dictionary that maps the source directory paths to the target directory paths.

Return type `dict`

groupby (*key=None, default=None*)

Group jobs according to one or more state point parameters.

This method can be called on any *JobsCursor* such as the one returned by *find_jobs()* or by iterating over a project.

Examples

```
# Group jobs by state point parameter 'a'.
for key, group in project.groupby('a'):
    print(key, list(group))

# Group jobs by document value 'a'.
for key, group in project.groupby('doc.a'):
    print(key, list(group))

# Group jobs by jobs.sp['a'] and job.document['b']
for key, group in project.groupby('a', 'doc.b'):
    print(key, list(group))

# Find jobs where job.sp['a'] is 1 and group them
# by job.sp['b'] and job.sp['c'].
for key, group in project.find_jobs({'a': 1}).groupby(('b', 'c')):
    print(key, list(group))

# Group by job.sp['d'] and job.document['count'] using a lambda.
for key, group in project.groupby(
    lambda job: (job.sp['d'], job.document['count'])
):
    print(key, list(group))
```

If *key* is *None*, jobs are grouped by id, placing one job into each group.

Parameters

- **key** (*str*, *iterable*, or *callable*) – The state point grouping parameter(s) passed as a string, iterable of strings, or a callable that will be passed one argument, the job (Default value = *None*).
- **default** – A default value to be used when a given state point key is not present. The value must be sortable and is only used if not *None* (Default value = *None*).

groupbydoc (*key=None, default=None*)

Group jobs according to one or more document values.

This method can be called on any *JobsCursor* such as the one returned by *find_jobs()* or by iterating over a project.

Examples

```
# Group jobs by document value 'a'.
for key, group in project.groupbydoc('a'):
    print(key, list(group))

# Find jobs where job.sp['a'] is 1 and group them
# by job.document['b'] and job.document['c'].
for key, group in project.find_jobs({'a': 1}).groupbydoc(('b', 'c')):
    print(key, list(group))
```

(continues on next page)

(continued from previous page)

```
# Group by whether 'd' is a field in the job.document using a lambda.
for key, group in project.groupbydoc(lambda doc: 'd' in doc):
    print(key, list(group))
```

If *key* is `None`, jobs are grouped by `id`, placing one job into each group.

Parameters

- **key** (*str, iterable, or callable*) – The document grouping parameter(s) passed as a string, iterable of strings, or a callable that will be passed one argument, *document* (Default value = `None`).
- **default** – A default value to be used when a given document key is not present. The value must be sortable and is only used if not `None` (Default value = `None`).

Deprecated since version 1.7: This will be removed in 2.0. Use `groupby` with a ‘`doc.`’ filter instead, see <https://docs.signac.io/en/latest/query.html#query-namespaces>.

`next()`

Return the next element.

This function is deprecated. Users should use `next(iter(...))` instead. .. deprecated:: 0.9.6

`to_dataframe(sp_prefix='sp.', doc_prefix='doc.', usecols=None, flatten=False)`

Convert the selection of jobs to a pandas `DataFrame`.

This function exports the job metadata to a `pandas.DataFrame`. All state point and document keys are prefixed by default to be able to distinguish them.

Parameters

- **sp_prefix** (*str, optional*) – Prefix state point keys with the given string. Defaults to “`sp.`”.
- **doc_prefix** (*str, optional*) – Prefix document keys with the given string. Defaults to “`doc.`”.
- **usecols** (*list-like or callable, optional*) – Used to select a subset of columns. If list-like, must contain strings corresponding to the column names that should be included. For example, [`'sp.a'`, `'doc.notes'`]. If callable, the column will be included if the function called on the column name returns `True`. For example, `lambda x: 'sp.' in x`. Defaults to `None`, which uses all columns from the state point and document. Note that this filter is applied *after* the `doc` and `sp` prefixes are added to the column names.
- **flatten** (*bool, optional*) – Whether nested state points or document keys should be flattened. If `True`, `{'a': {'b': 'c'}}` becomes a column named `a.b` with value `c`. If `False`, it becomes a column named `a` with value `{'b': 'c'}`. Defaults to `False`.

Returns A pandas `DataFrame` with all job metadata.

Return type `DataFrame`

1.3.3 The Job class

Attributes

<code>Job.clear()</code>	Remove all job data, but not the job itself.
<code>Job.close()</code>	Close the job and switch to the previous working directory.
<code>Job.data</code>	Get data associated with this job.
<code>Job.doc</code>	Alias for <code>document</code> .
<code>Job.document</code>	Get document associated with this job.
<code>Job.fn(filename)</code>	Prepend a filename with the job's workspace directory path.
<code>Job.get_id()</code>	Job's state point unique identifier.
<code>Job.id</code>	Get the unique identifier for the job's state point.
<code>Job.init([force])</code>	Initialize the job's workspace directory.
<code>Job.isfile(filename)</code>	Return True if file exists in the job's workspace.
<code>Job.move(project)</code>	Move this job to project.
<code>Job.open()</code>	Enter the job's workspace directory.
<code>Job.remove()</code>	Remove the job's workspace including the job document.
<code>Job.reset()</code>	Remove all job data, but not the job itself.
<code>Job.reset_statepoint(new_statepoint)</code>	Overwrite the state point of this job while preserving job data.
<code>Job.sp</code>	Alias for <code>statepoint</code> .
<code>Job.statepoint</code>	Get the job's state point.
<code>Job.stores</code>	Get HDF5 stores associated with this job.
<code>Job.sync(other[, strategy, exclude, doc_sync])</code>	Perform a one-way synchronization of this job with the other job.
<code>Job.update_statepoint(update[, overwrite])</code>	Change the state point of this job while preserving job data.
<code>Job.workspace()</code>	Return the job's unique workspace directory.
<code>Job.ws</code>	Alias for <code>workspace()</code> .

class `signac.contrib.job.Job` (*project*, *statepoint=None*, *_id=None*)

Bases: `object`

The job instance is a handle to the data of a unique state point.

Application developers should not directly instantiate this class, but use `open_job()` instead.

Jobs can be opened by `statepoint` or `_id`. If both values are provided, it is the user's responsibility to ensure that the values correspond.

Parameters

- **project** (*Project*) – Project handle.
- **statepoint** (*dict*) – State point for the job. (Default value = None)
- **_id** (*str*) – The job identifier. (Default value = None)

FN_DOCUMENT = `'signac_job_document.json'`

The job's document filename.

FN_MANIFEST = `'signac_statepoint.json'`

The job's state point filename.

The job state point is a human-readable file containing the job's state point that is stored in each job's workspace directory.

KEY_DATA = `'signac_data'`

The job's datastore key.

clear()

Remove all job data, but not the job itself.

This function will do nothing if the job was not previously initialized.

See *signac rm -c* for the command line equivalent.

close()

Close the job and switch to the previous working directory.

data

Get data associated with this job.

This property should be used for large array-like data, which can't be stored efficiently in the job document. For examples and usage, see [Job Data Storage](#).

Equivalent to:

```
return job.stores['signac_data']
```

Returns An HDF5-backed datastore.

Return type *H5Store*

doc

Alias for *document*.

Warning: Even deep copies of *doc* will modify the same file, so changes will still effectively be persisted between deep copies. If you need a deep copy that will not modify the underlying persistent JSON file, use the call operator to get an equivalent plain dictionary: `job.doc()`.

See *signac document* for the command line equivalent.

Returns The job document handle.

Return type *JSONDict*

document

Get document associated with this job.

Warning: Even deep copies of *document* will modify the same file, so changes will still effectively be persisted between deep copies. If you need a deep copy that will not modify the underlying persistent JSON file, use the call operator to get an equivalent plain dictionary: `job.document()`. For more information, see *JSONDict*.

See *signac document* for the command line equivalent.

Returns The job document handle.

Return type *JSONDict*

fn(filename)

Prepend a filename with the job's workspace directory path.

Parameters **filename** (*str*) – The name of the file.

Returns The full workspace path of the file.

Return type *str*

get_id()

Job's state point unique identifier.

Returns The job id.

Return type `str`

Deprecated since version 1.3: This will be removed in 2.0. Use `job.id` instead.

id

Get the unique identifier for the job's state point.

Returns The job id.

Return type `str`

init (*force=False*)

Initialize the job's workspace directory.

This function will do nothing if the directory and the job state point already exist and the state point is valid.

Returns the calling job.

See `signac job -c` for the command line equivalent.

Parameters **force** (*bool*) – Overwrite any existing state point files, e.g., to repair them if they got corrupted (Default value = False).

Returns The job handle.

Return type `Job`

Raises

- `OSError` – If the workspace directory cannot be created or any other I/O error occurs when attempting to save the state point file.
- `JobsCorruptedError` – If the job state point on disk is corrupted.

isfile (*filename*)

Return True if file exists in the job's workspace.

Parameters **filename** (*str*) – The name of the file.

Returns True if file with filename exists in workspace.

Return type `bool`

move (*project*)

Move this job to project.

This function will attempt to move this instance of job from its original project to a different project.

See `signac move` for the command line equivalent.

Parameters **project** (*Project*) – The project to move this job to.

open ()

Enter the job's workspace directory.

You can use the `Job` class as context manager:

```
with project.open_job(my_statepoint) as job:
    # manipulate your job data
```

Opening the context will switch into the job’s workspace, leaving it will switch back to the previous working directory.

remove ()

Remove the job’s workspace including the job document.

This function will do nothing if the workspace directory does not exist.

See *signac rm* for the command line equivalent.

reset ()

Remove all job data, but not the job itself.

This function will initialize the job if it was not previously initialized.

reset_statepoint (new_statepoint)

Overwrite the state point of this job while preserving job data.

This method will change the job id if the state point has been altered.

For more information, see [Modifying the State Point](#).

Danger: Use this function with caution! Resetting a job’s state point may sometimes be necessary, but can possibly lead to incoherent data spaces.

Parameters `new_statepoint (dict)` – The job’s new state point.

sp

Alias for *statepoint*.

statepoint

Get the job’s state point.

Warning: The state point object behaves like a dictionary in most cases, but because it persists changes to the filesystem, making a copy requires explicitly converting it to a dict. If you need a modifiable copy that will not modify the underlying JSON file, you can access a dict copy of the state point by calling it, e.g. `sp_dict = job.statepoint ()` instead of `sp = job.statepoint`. For more information, see *JSONAttrDict*.

See *signac statepoint* for the command line equivalent.

Returns Returns the job’s state point.

Return type `dict`

stores

Get HDF5 stores associated with this job.

Use this property to access an HDF5 file within the job’s workspace directory using the *H5Store* dict-like interface.

This is an example for accessing an HDF5 file called ‘my_data.h5’ within the job’s workspace:

```
job.stores['my_data']['array'] = np.random((32, 4))
```

This is equivalent to:

```
H5Store(job.fn('my_data.h5'))['array'] = np.random((32, 4))
```

Both the `stores` and the `H5Store` itself support attribute access. The above example could therefore also be expressed as:

```
job.stores.my_data.array = np.random((32, 4))
```

Returns The HDF5-Store manager for this job.

Return type `H5StoreManager`

sync (*other*, *strategy=None*, *exclude=None*, *doc_sync=None*, ***kwargs*)

Perform a one-way synchronization of this job with the other job.

By default, this method will synchronize all files and document data with the other job to this job until a synchronization conflict occurs. There are two different kinds of synchronization conflicts:

1. The two jobs have files with the same, but different content.
2. The two jobs have documents that share keys, but those keys are associated with different values.

A file conflict can be resolved by providing a ‘FileSync’ *strategy* or by *excluding* files from the synchronization. An unresolvable conflict is indicated with the raise of a `FileSyncConflict` exception.

A document synchronization conflict can be resolved by providing a `doc_sync` function that takes the source and the destination document as first and second argument.

Parameters

- **other** (`Job`) – The other job to synchronize from.
- **strategy** – A synchronization strategy for file conflicts. If no strategy is provided, a `SyncConflict` exception will be raised upon conflict (Default value = None).
- **exclude** (`str`) – An filename exclude pattern. All files matching this pattern will be excluded from synchronization (Default value = None).
- **doc_sync** – A synchronization strategy for document keys. If this argument is None, by default no keys will be synchronized upon conflict.
- **dry_run** – If True, do not actually perform the synchronization.
- ****kwargs** – Extra keyword arguments will be forward to the `sync_jobs()` function which actually excutes the synchronization operation.

Raises `FileSyncConflict` – In case that a file synchronization results in a conflict.

update_statepoint (*update*, *overwrite=False*)

Change the state point of this job while preserving job data.

By default, this method will not change existing parameters of the state point of the job.

This method will change the job id if the state point has been altered.

For more information, see [Modifying the State Point](#).

Warning: While appending to a job’s state point is generally safe, modifying existing parameters may lead to data inconsistency. Use the `overwrite` argument with caution!

Parameters

- **update** (`dict`) – A mapping used for the state point update.

- **overwrite** (*bool, optional*) – If False, an error will be raised if the update modifies the values of existing keys in the state point. If True, any existing keys will be overwritten in the same way as `dict.update()`. Use with caution! (Default value = False).

Raises

- `KeyError` – If the update contains keys which are already part of the job’s state point and `overwrite` is False.
- `DestinationExistsError` – If a job associated with the new state point is already initialized.
- `OSError` – If the move failed due to an unknown system related error.

workspace()

Return the job’s unique workspace directory.

See `signac job -w` for the command line equivalent.

Returns The path to the job’s workspace directory.

Return type `str`

ws

Alias for `workspace()`.

1.3.4 The Collection

class `signac.Collection` (*docs=None, primary_key='_id', compresslevel=0, _trust=False*)

A collection of documents.

The Collection class manages a collection of documents in memory or in a file on disk. A document is defined as a dictionary mapping of key-value pairs.

An instance of collection may be used to manage and search documents. For example, given a collection with member data, where each document contains a *name* entry and an *age* entry, we can find the name of all members that are at age 32 like this:

```
members = [
    {'name': 'John', 'age': 32},
    {'name': 'Alice', 'age': 28},
    {'name': 'Kevin', 'age': 32},
    # ...
]

member_collection = Collection(members)
for doc in member_collection.find({'age': 32}):
    print(doc['name'])
```

To iterate over all documents in the collection, use:

```
for doc in collection:
    print(doc)
```

By default a collection object will reside in memory. However, it is possible to manage a collection associated to a file on disk. To open a collection which is associated with a file on disk, use the `Collection.open()` class method:

```
with Collection.open('collection.txt') as collection:
    for doc in collection.find({'age': 32}):
        print(doc)
```

The collection file is by default opened in *a+* mode, which means it can be read from and written to and will be created if it does not exist yet.

Parameters

- **docs** (*iterable*) – Initialize the collection with these documents.
- **primary_key** (*str*) – The name of the key which serves as the primary index of the collection. Selecting documents by primary key has time complexity of $O(N)$ in the worst case and $O(1)$ on average. All documents must have a primary key value. The default primary key is *_id*.
- **compresslevel** (*int*) – The level of compression to use. Any positive value implies compression and is used by the underlying *gzip* implementation. Default value is 0 (no compression).

Raises `ValueError` – When first argument is a string.

`clear()`

Remove all documents from the collection.

`close()`

Close this collection instance.

In case that the collection is associated with a file-object, all changes are flushed to the file and the file is closed.

It is not possible to re-open the same collection instance after closing it.

`delete_many(filter)`

Delete all documents that match the filter.

Parameters **filter** (*dict*) – A document that should be deleted must match this filter.

`delete_one(filter)`

Delete one document that matches the filter.

Parameters **filter** (*dict*) – The document that should be deleted must match this filter.

`dump(file=<io.TextIOWrapper name='<stdout>' mode='w' encoding='UTF-8'>)`

Dump the collection in JSON-encoding to file.

The file argument defaults to *sys.stdout*, which means the encoded blob will be printed to screen in case that no file argument is provided.

For example, to dump to a file on disk, one could write:

```
with open('my_collection.txt', 'w') as file:
    collection.dump(file)
```

Parameters **file** – The file to write the encoded blob to (Default value = *sys.stdout*).

`find(filter=None, limit=0)`

Find all documents matching filter, but not more than limit.

This function searches the collection for all documents that match the given filter and returns a result vector. For example:

```
for doc in collection.find(my_filter):
    print(doc)
```

Nested values should be searched using the `.` operator, for example:

```
docs = collection.find({'nested.value': 42})
```

will return documents with a nested structure: `{'nested': {'value': 42}}`.

The result of `find()` can be stored and iterated over multiple times. In addition, the result vector can be queried for its size:

```
docs = collection.find(my_filter)

print(len(docs))      # the number of documents matching

for doc in docs:      # iterate over the result vector
    pass
```

Arithmetic Operators

- `$eq`: equal
- `$ne`: not equal
- `$gt`: greater than
- `$gte`: greater or equal than
- `$lt`: less than
- `$lte`: less or equal than

```
project.find({"a": {"$lt": 5}})
```

Matches all docs with `a` less than 5.

Logical Operators

That includes `$and` and `$or`; both expect a list of expressions.

```
project.find({"$or": [{"a": 4}, {"b": {"$gt": 3}}]})
```

Matches all docs, where `a` is 4 or `b` is greater than 3.

Exists operator

Determines whether a specific key exists, or not, e.g.:

```
project.find({"a": {"$exists": True}})
```

Array operator

To determine whether specific elements are in (`$in`), or not in (`$nin`) an array, e.g.:

```
project.find({"a": {"$in": [0, 1, 2]}})
```

Matches all docs, where `a` is either 0, 1, or 2. Usage of `$nin` is equivalent.

Regular expression operator

Allows the “on-the-fly” evaluation of regular expressions, e.g.:

```
project.find({"protocol": {"$regex": "foo"}})
```

Will match all docs with a protocol that contains the term 'foo'.

\$type operator

Matches when a value is of specific type, e.g.:

```
project.find({"protocol": {"$type": str}})
```

Finds all docs, where the value of protocol is of type str. Other types that can be checked are: *int*, *float*, *bool*, *list*, and *null*.

\$where operator

Matches an arbitrary python expression, e.g.:

```
project.find({"foo": {"$where": "lambda x: x.startswith('bar')"}}
↪)
```

Matches all docs, where the value for foo starts with the word 'bar'.

Parameters

- **filter** (*dict*) – All documents must match the given filter (Default value = None).
- **limit** (*int*) – Do not return more than limit number of documents. A limit value of 0 (the default) means no limit.

Returns A result object that iterates over all matching documents.

Return type `_CollectionSearchResults`

Raises `ValueError` – In case that the filter argument is invalid.

find_one (*filter=None*)

Return one document that matches the filter or None.

```
doc = collection.find_one(my_filter)
if doc is None:
    print("No result found for filter", my_filter)
else:
    print("Doc matching filter:", my_filter, doc)
```

Parameters **filter** (*dict*) – The returned document must match the given filter (Default value = None).

Returns A matching document or None.

Return type `dict`

Raises `ValueError` – In case that the filter argument is invalid.

flush ()

Write all changes to the associated file.

If the collection instance is associated with a file-object, calling the `flush()` method will write all changes to this file.

This method is also called when the collection is explicitly or implicitly closed.

ids

Get an iterator over the primary key in the collection.

Returns iterator over the primary key in the collection.

Return type iterable

index (*key*, *build=False*)

Get (and optionally build) the index for a given key.

An index allows to access documents by a specific key with minimal time complexity, e.g.:

```
age_index = member_collection.index('age')
for _id in age_index[32]:
    print(member_collection[_id]['name'])
```

This means we can access documents by the 'age' key in O(1) time on average in addition to the primary key. Using the *find()* method will automatically build all required indexes for the particular search.

Once an index has been built, it will be internally managed by the class and updated with subsequent changes. An index returned by this method is always current with the latest state of the collection.

Parameters

- **key** (*str*) – The primary key of the requested index.
- **build** (*bool*) – If True, build a non-existing index if necessary, otherwise raise `KeyError` (Default value = False).

Returns Index for the given key.

Return type dict

Raises `KeyError` – In case the build is False and the index has not been built yet or no index is present for the key.

insert_one (*doc*)

Insert one document into the collection.

If the document does not have a value for the collection's primary key yet, it will be assigned one.

```
_id = collection.insert_one(doc)
assert _id in collection
```

Note: The document will be directly updated in case that it has no primary key and must therefore be mutable!

Parameters **doc** (*dict*) – The document to be inserted.

Returns The *_id* of the inserted document.

Return type str

main ()

Start a command line interface for this Collection.

Use this function to interact with this instance of Collection on the command line. For example, executing the following script:

```
# find.py
with Collection.open('my_collection.txt') as c:
    c.main()
```

will enable us to search for documents on the command line like this:

```
$ python find.py '{"age": 32}'
{"name": "John", "age": 32}
{"name": "Kevin", "age": 32}
```

Raises `ValueError` – When both `-id` or `-indent` are selected.

classmethod `open` (*filename*, *mode=None*, *compresslevel=None*)

Open a collection associated with a file on disk.

Using this factory method will return a collection that is associated with a collection file on disk. For example:

```
with Collection.open('collection.txt') as collection:
    for doc in collection:
        print(doc)
```

will read all documents from the `collection.txt` file or create the file if it does not exist yet.

Modifications to the file will be written to the file when the `flush()` method is called or the collection is explicitly closed by calling the `Collection.close()` method or implicitly by leaving the `with`-clause:

```
with Collection.open('collection.txt') as collection:
    collection.update(my_docs)
# All changes to the collection have been written to collection.txt.
```

The open-modes work as expected, so for example to open a collection file in *read-only* mode, use `Collection.open('collection.txt', 'r')`.

Opening a gzip (`*.gz`) file also works as expected. Because gzip does not support a combined read and write mode, `mode=*` is not available. Be sure to open the file in read, write, or append mode as required. Due to the manner in which gzip works, opening a file in `mode=wt` will effectively erase the current file, so take care using `mode=wt`.

Parameters

- **filename** (*str*) – Name of file to read the documents from or create the file if it does not exist.
- **mode** (*str*) – Open the file with mode (Default value = None).
- **compresslevel** (*int*) – The level of compression to use. Any positive value implies compression and is used by the underlying gzip implementation. (Default value = None)

Returns An instance of `Collection`.

Return type `Collection`

Raises `RuntimeError` – File open-mode is not None for in-memory collection or compressed collections are not opened in binary mode.

primary_key

Get the name of the collection's primary key (default='_id').

classmethod `read_json` (*file=None*)

Construct an instance of `Collection` from a JSON file.

Parameters **file** – The json file to read, provided as either a filename or a file-like object (Default value = None).

Returns A Collection containing the JSON file

Return type *Collection*

replace_one (*filter, replacement, upsert=False*)

Replace one document that matches the given filter.

The first document matching the filter will be replaced by the given replacement document. If the *upsert* argument is True, the replacement will be inserted in case that no document matches the filter.

Parameters

- **filter** (*dict*) – A document that should be replaced must match this filter.
- **replacement** (*dict*) – The replacement document.
- **upsert** (*bool*) – If True, insert the replacement document in the case that no document matches the filter (Default value = False).

Returns The id of the replaced (or upserted) documented.

Return type *str*

Raises *ValueError* – In case that the filter argument is invalid.

to_json (*file=None*)

Dump the collection as a JSON file.

This function returns the JSON-string directly if the file argument is None.

Parameters **file** – The filename or a file-like object to write the JSON string to (Default value = None).

Returns JSON-string when file argument is not provided.

Return type *JSON*

update (*docs*)

Update the collection with these documents.

Any existing documents with the same primary key will be replaced.

Parameters **docs** (*iterable*) – A sequence of documents to be upserted into the collection.

1.3.5 The JSONDict

This class implements the interface for the job's *statepoint* and *document* attributes, but can also be used stand-alone:

```
signac.JSONDict
    alias      of      signac.synced_collections.backends.collection_json.
                  BufferedJSONAttrDict
```

1.3.6 The H5Store

This class implements the interface to the job's *data* attribute, but can also be used stand-alone:

class `signac.H5Store` (*filename*, ***kwargs*)

An HDF5-backed container for storing array-like and dictionary-like data.

The H5Store is a `MutableMapping` and therefore behaves similar to a `dict`, but all data is stored persistently in the associated HDF5 file on disk.

Supported types include:

- built-in types (int, float, str, bool, NoneType, array)
- numpy arrays
- pandas data frames (requires pandas and pytables)
- mappings with values that are supported types

Values can be accessed as attributes (`h5s.foo`) or via key index (`h5s['foo']`).

Examples

```
>>> from signac import H5Store
>>> with H5Store('file.h5') as h5s:
...     h5s['foo'] = 'bar'
...     assert 'foo' in h5s
...     assert h5s.foo == 'bar'
...     assert h5s['foo'] == 'bar'
>>>
```

The H5Store can be used as a context manager to ensure that the underlying file is opened, however most built-in types (excluding arrays) can be read and stored without the need to `_explicitly_` open the file. **To access arrays (reading or writing), the file must always be opened!**

To open a file in read-only mode, use the `open()` method with `mode='r'`:

```
>>> with H5Store('file.h5').open(mode='r') as h5s:
...     pass
>>>
```

Parameters

- **filename** (*str*) – The filename of the underlying HDF5 file.
- ****kwargs** – Additional keyword arguments to be forwarded to the `h5py.File` constructor. See the documentation for the `h5py.File` constructor for more information.

clear()

Remove all data from this store.

Danger: All data will be removed, this action cannot be reversed!

close()

Close the underlying HDF5 file.

file

Access the underlying instance of `h5py.File`.

This property exposes the underlying `h5py.File` object enabling use of functions such as `create_dataset()` or `requires_dataset()`.

Note: The store must be open to access this property!

Returns The `h5py` file-object that this store is operating on.

Return type `h5py.File`

Raises **H5StoreClosedError** – When the store is closed at the time of accessing this property.

filename

Return the `H5Store` filename.

flush()

Flush the underlying HDF5 file.

get (k , d) → $D[k]$ if k in D , else d . d defaults to `None`.

items() → a set-like object providing a view on D 's items

keys() → a set-like object providing a view on D 's keys

mode

Return the default opening mode of this `H5Store`.

open ($mode=None$)

Open the underlying HDF5 file.

Parameters **mode** – The file open mode to use. Defaults to 'a' (append).

Returns This `H5Store` instance.

pop (k , d) → v , remove specified key and return the corresponding value.
If key is not found, d is returned if given, otherwise `KeyError` is raised.

popitem() → (k , v), remove and return some (key, value) pair
as a 2-tuple; but raise `KeyError` if D is empty.

setdefault (key , $value$)

Set a value for a key if that key is not already set.

update ($[E]$, $**F$) → `None`. Update D from mapping/iterable E and F .

If E present and has a `.keys()` method, does: for k in E : $D[k] = E[k]$ If E present and lacks `.keys()` method, does: for (k , v) in E : $D[k] = v$ In either case, this is followed by: for k , v in $F.items()$: $D[k] = v$

values() → an object providing a view on D 's values

1.3.7 The `H5StoreManager`

This class implements the interface to the job's `stores` attribute, but can also be used stand-alone:

```
class signac.H5StoreManager (prefix)
```

```
    Bases: signac.core.dict_manager.DictManager
```

```
    Helper class to manage multiple instances of H5Store within a directory.
```

```
    Example (assuming that the 'stores/' directory exists):
```

```

>>> stores = H5StoreManager('stores/')
>>> stores.data
<H5Store(filename=stores/data.h5)>
>>> stores.data.foo = True
>>> dict(stores.data)
{'foo': True}

```

Parameters `prefix` – The directory prefix shared by all stores managed by this class.

cls

alias of `H5Store`

keys()

Return an iterable of keys.

prefix

Return the prefix.

1.3.8 Top-level functions

The signac framework aids in the management of large and heterogeneous data spaces.

It provides a simple and robust data model to create a well-defined, indexable storage layout for data and metadata. This makes it easier to operate on large data spaces, streamlines post-processing and analysis, and makes data collectively accessible.

`signac.TemporaryProject` (*name=None, cls=None, **kwargs*)

Context manager for the generation of a temporary project.

This is a factory function that creates a `Project` within a temporary directory and must be used as context manager, for example like this:

```

with TemporaryProject() as tmp_project:
    tmp_project.import_from('/data')

```

Parameters

- **name** (*str*) – An optional name for the temporary project. Defaults to a unique random string.
- **cls** – The class of the temporary project. Defaults to `Project`.
- ****kwargs** – Optional keyword arguments that are forwarded to the `TemporaryDirectory` class constructor, which is used to create a temporary root directory.

Yields `Project` – An instance of `Project`.

`signac.get_project` (*root=None, search=True, **kwargs*)

Find a project configuration and return the associated project.

Parameters

- **root** (*str*) – The starting point to search for a project, defaults to the current working directory.
- **search** (*bool*) – If `True`, search for project configurations inside and above the specified root directory, otherwise only return projects with a root directory identical to the specified root argument (Default value = `True`).

- ****kwargs** – Optional keyword arguments that are forwarded to `get_project()`.

Returns An instance of `Project`.

Return type `Project`

Raises `LookupError` – If no project configuration can be found.

`signac.init_project` (*name*, *root=None*, *workspace=None*, *make_dir=True*)

Initialize a project with the given name.

It is safe to call this function multiple times with the same arguments. However, a `RuntimeError` is raised if an existing project configuration would conflict with the provided initialization parameters.

Parameters

- **name** (*str*) – The name of the project to initialize.
- **root** (*str*) – The root directory for the project. Defaults to the current working directory.
- **workspace** (*str*) – The workspace directory for the project. Defaults to a subdirectory `workspace` in the project root.
- **make_dir** (*bool*) – Create the project root directory, if it does not exist yet (Default value = `True`).

Returns The initialized project instance.

Return type `Project`

Raises `RuntimeError` – If the project root path already contains a conflicting project configuration.

`signac.get_job` (*root=None*)

Find a Job in or above the current working directory (or provided path).

Parameters **root** (*str*) – The job root directory. If no root directory is given, the current working directory is assumed to be within the current job workspace directory (Default value = `None`).

Returns Job handle.

Return type `Job`

Raises `LookupError` – If this job cannot be found.

Examples

When the current directory is a job workspace directory:

```
>>> signac.get_job()
signac.contrib.job.Job(project=..., statepoint={...})
```

`signac.diff_jobs` (**jobs*)

Find differences among a list of jobs' state points.

The resulting diff is a dictionary where the keys are job ids and the values are each job's state point minus the intersection of all provided jobs' state points. The comparison is performed over the combined set of keys and values.

See `signac diff` for the command line equivalent.

Parameters ***jobs** (sequence[`Job`]) – Sequence of jobs to diff.

Returns A dictionary where the keys are job ids and values are the unique parts of that job's state point.

Return type dict

Examples

```
>>> import signac
>>> project = signac.init_project('project_name')
>>> job1 = project.open_job({'constant': 42, 'diff1': 0, 'diff2': 1}).init()
>>> job2 = project.open_job({'constant': 42, 'diff1': 1, 'diff2': 1}).init()
>>> job3 = project.open_job({'constant': 42, 'diff1': 2, 'diff2': 2}).init()
>>> print(job1)
c4af2b26f1fd256d70799ad3ce3bdad0
>>> print(job2)
b96b21fada698f8934d58359c72755c0
>>> print(job3)
e4289419d2b0e57e4852d44a09f167c0
>>> signac.diff_jobs(job1, job2, job3)
{'c4af2b26f1fd256d70799ad3ce3bdad0': {'diff2': 1, 'diff1': 0},
 'b96b21fada698f8934d58359c72755c0': {'diff2': 1, 'diff1': 1},
 'e4289419d2b0e57e4852d44a09f167c0': {'diff2': 2, 'diff1': 2}}
>>> signac.diff_jobs(*project)
{'c4af2b26f1fd256d70799ad3ce3bdad0': {'diff2': 1, 'diff1': 0},
 'b96b21fada698f8934d58359c72755c0': {'diff2': 1, 'diff1': 1},
 'e4289419d2b0e57e4852d44a09f167c0': {'diff2': 2, 'diff1': 2}}
```

`signac.get_database(name, hostname=None, config=None)`

Get a database handle.

The database handle is an instance of `Database`, which provides access to the document collections within one database.

```
db = signac.db.get_database('MyDatabase')
docs = db.my_collection.find()
```

Please note, that a collection which did not exist at the point of access, will automatically be created.

Parameters

- **name** (*str*) – The name of the database to get.
- **hostname** (*str*) – The name of the configured host. Defaults to the first configured host, or the host specified by `default_host`.
- **config** (`common.config.Config`) – The config object to retrieve the host configuration from. Defaults to the global configuration.

Returns The database handle.

Return type `pymongo.database.Database`

See also:

<https://api.mongodb.org/python/current/api/pymongo/database.html>

Deprecated since version 1.3: This will be removed in 2.0. The database package is deprecated.

`signac.fetch(doc_or_id, mode='r', mirrors=None, num_tries=3, timeout=60, ignore_local=False)`

Fetch the file associated with this document or file id.

This function retrieves a file associated with the provided index document or file id and behaves like the built-in `open()` function, e.g.:

```
for doc in index:
    with signac.fetch(doc) as file:
        do_something_with(file)
```

Parameters

- **doc_or_id** – A file_id or a document with a file_id value.
- **mode** – Mode to use for opening files.
- **mirrors** – An optional set of mirrors to fetch the file from.
- **num_tries** (*int*) – The number of automatic retry attempts in case of mirror connection errors.
- **timeout** (*int*) – The time in seconds to wait before an automatic retry attempt.

Returns The file associated with the document or file id.

Return type A file-like object

Deprecated since version 1.3: This will be removed in 2.0. The indexing module is deprecated.

`signac.export_one(doc, index, mirrors=None, num_tries=3, timeout=60)`

Export one document to index and an optionally associated file to mirrors.

Parameters

- **doc** – A document with a file_id entry.
- **docs** – The index collection to export to.
- **mirrors** – An optional set of mirrors to export files to.
- **num_tries** (*int*) – The number of automatic retry attempts in case of mirror connection errors.
- **timeout** (*int*) – The time in seconds to wait before an automatic retry attempt.

Returns The id and file id after successful export.

Deprecated since version 1.3: This will be removed in 2.0. The indexing module is deprecated.

`signac.export(docs, index, mirrors=None, update=False, num_tries=3, timeout=60, **kwargs)`

Export docs to index and optionally associated files to mirrors.

The behavior of this function is equivalent to:

```
for doc in docs:
    export_one(doc, index, mirrors, num_tries)
```

If the *update* argument is set to True, the export algorithm will automatically identify stale index documents, that means documents that refer to files or state points that have been removed and are no longer part of the data space. Any document which shares the *root*, but not the *_id* field with any of the updated documents is considered stale and removed. Using *update* in combination with an empty docs sequence will raise *ExportError*, since it is not possible to identify stale documents in that case.

Note: This function will automatically delegate to specialized implementations for special index types. For example, if the index argument is a MongoDB document collection, the index documents will be exported via `export_pymongo()`.

Parameters

- **docs** – The index documents to export.
- **index** – The collection to export the index to.
- **mirrors** – An optional set of mirrors to export files to.
- **update** (*bool*) – If True, remove stale index documents, that means documents that refer to files or state points that no longer exist.
- **num_tries** (*int*) – The number of automatic retry attempts in case of mirror connection errors.
- **timeout** (*int*) – The time in seconds to wait before an automatic retry attempt.
- **kwargs** – Optional keyword arguments to pass to delegate implementations.

Raises `ExportError` – When using the update argument in combination with an empty docs sequence.

Deprecated since version 1.3: This will be removed in 2.0. The indexing module is deprecated.

`signac.export_to_mirror` (*doc*, *mirror*, *num_tries*=3, *timeout*=60)

Export a file associated with doc to mirror.

Parameters

- **doc** – A document with a `file_id` entry.
- **mirror** – A file-system object to export the file to.
- **num_tries** (*int*) – The number of automatic retry attempts in case of mirror connection errors.
- **timeout** (*int*) – The time in seconds to wait before an automatic retry attempt.

Returns The file id after successful export.

Deprecated since version 1.3: This will be removed in 2.0. The indexing module is deprecated.

`signac.export_pymongo` (*docs*, *index*, *mirrors*=None, *update*=False, *num_tries*=3, *timeout*=60, *chunk-size*=100)

Optimized `export()` function for pymongo index collections.

The behavior of this function is roughly equivalent to:

```
for doc in docs:
    export_one(doc, index, mirrors, num_tries)
```

Note: All index documents must be JSON-serializable to be able to be exported to a MongoDB collection.

Parameters

- **docs** – The index documents to export.
- **index** (`pymongo.collection.Collection`) – The database collection to export the index to.
- **num_tries** (*int*) – The number of automatic retry attempts in case of mirror connection errors.
- **timeout** (*int*) – The time in seconds to wait before an automatic retry attempt.

- **chunksize** (*int*) – The buffer size for export operations.

Deprecated since version 1.3: This will be removed in 2.0. The indexing module is deprecated.

`signac.index_files` (*root='.', formats=None, depth=0*)

Generate a file index.

This generator function yields file index documents, where each index document corresponds to one file.

To index all files in the current working directory, simply execute:

```
for doc in signac.index_files():
    print(doc)
```

A file associated with a file index document can be fetched via the `fetch()` function:

```
for doc in signac.index_files():
    with signac.fetch(doc) as file:
        print(file.read())
```

This is especially useful if the file index is part of a collection (*Collection*) which can be searched for specific entries.

To limit the file index to files with a specific filename formats, provide a regular expression as the `formats` argument. To index all files that have file ending `.txt`, execute:

```
for doc in signac.index_files(formats='.*\.txt'):
    print(doc)
```

We can specify specific formats by providing a dictionary as `formats` argument, where the key is the filename pattern and the value is an arbitrary formats string, e.g.:

```
for doc in signac.index_files(formats={
    r'.*\.txt': 'TextFile', r'.*\.zip': 'ZipFile'}):
    print(doc)
```

Parameters

- **root** (*str*) – The directory to index, defaults to the current working directory.
- **formats** – Limit the index to files that match the given regular expression and optionally associate formats with given patterns.
- **depth** (*int*) – Limit the search to the specified directory depth.

Yields The file index documents as dicts.

Deprecated since version 1.3: This will be removed in 2.0. The indexing module is deprecated.

`signac.index` (*root='.', tags=None, depth=0, **kwargs*)

Generate a main index.

A main index is compiled from other indexes by searching for modules named `signac_access.py` and compiling all indexes which are yielded from a function `get_indexes(root)` defined within that module as well as the indexes generated by crawlers yielded from a function `get_crawlers(root)` defined within that module.

This is a minimal example for a `signac_access.py` file:

```
import signac

def get_indexes(root):
    yield signac.index_files(root, r'.*\.txt')
```

Internally, this function constructs an instance of `MainCrawler` and all extra key-word arguments will be forwarded to the constructor of said main crawler.

Parameters

- **root** (*str*) – Look for access modules under this directory path.
- **tags** – If tags are provided, do not execute subcrawlers that don't match the same tags.
- **depth** (*int*) – Limit the search to the specified directory depth.
- **kwargs** – These keyword-arguments are forwarded to the internal `MainCrawler` instance.

Yields The main index documents as instances of `dict`.

Deprecated since version 1.3: This will be removed in 2.0. The indexing module is deprecated.

`signac.flush()`

Execute all deferred `JSONDict` write operations.

Deprecated since version 1.7: This will be removed in 2.0.

1.3.9 Submodules

signac.cite module

Functions to support citing this software.

`signac.cite.bibtex` (*file=None*)

Generate bibtex entries for signac.

The bibtex entries will be printed to screen unless a filename or a file-like object are provided, in which case they will be written to the corresponding file.

Note: A full reference should also include the version of this software. Please refer to the documentation on how to cite a specific version.

Parameters **file** – A str or file-like object. Defaults to `sys.stdout`.

`signac.cite.reference` (*file=None*)

Generate formatted reference entries for signac.

The references will be printed to screen unless a filename or a file-like object are provided, in which case they will be written to the corresponding file.

Note: A full reference should also include the version of this software. Please refer to the documentation on how to cite a specific version.

Parameters **file** – A str or file-like object. Defaults to `sys.stdout`.

signac.sync module

Synchronization of jobs and projects.

Jobs may be synchronized by copying all data from the source job to the destination job. This means all files are copied and the documents are synchronized. Conflicts, that means both jobs contain conflicting data, may be resolved with a user defined strategy.

The synchronization of projects is in essence the synchronization of all jobs which are in the destination project with the ones in the source project and the sync synchronization of the project document. If a specific job does not exist yet at the destination it is simply cloned, otherwise it is synchronized.

A sync strategy is a function (or functor) that takes the source job, the destination job, and the name of the file generating the conflict as arguments and returns the decision whether to overwrite the file as Boolean. There are some default strategies defined within this module as part of the `FileSync` class. These are the default strategies:

1. `always` – Always overwrite on conflict.
2. `never` – Never overwrite on conflict.
3. `update` – Overwrite when the modification time of the source file is newer.
4. `Ask` – Ask the user interactively about each conflicting filename.

For example, to synchronize two projects resolving conflicts by modification time, use:

```
dest_project.sync(source_project, strategy=sync.FileSync.update)
```

Unlike files, which are always either overwritten as a whole or not, documents can be synchronized more fine-grained with a *sync function*. Such a function (or functor) takes the source and the destination document as arguments and performs the synchronization. The user is encouraged to implement their own sync functions, but there are a few default functions implemented as part of the `DocSync` class:

1. `NO_SYNC` – Do not perform any synchronization.
2. `COPY` – Apply the same strategy used to resolve file conflicts.
3. `update` – Equivalent to `dst.update(src)`.
4. `ByKey` – Synchronize the source document key by key, more information below.

This is how we could synchronize two jobs, where the documents are synchronized with a simple update function:

```
dst_job.sync(src_job, doc_sync=sync.DocSync.update)
```

The `DocSync.ByKey` functor attempts to synchronize the destination document with the source document without overwriting any data. That means this function behaves similar to `update()` for a non-intersecting set of keys, but in addition will preserve nested mappings without overwriting values. In addition, any key conflict, that means keys that are present in both documents, but have differing data, will lead to the raise of a `DocumentSyncConflict` exception. The user may explicitly decide to overwrite certain keys by providing a “key-strategy”, which is a function that takes the conflicting key as argument, and returns the decision whether to overwrite that specific key as Boolean. For example, to sync two jobs, where conflicting keys should only be overwritten if they contain the term ‘foo’, we could execute:

```
dst_job.sync(src_job, doc_sync=sync.DocSync.ByKey(lambda key: 'foo' in key))
```

This means that all documents are synchronized ‘key-by-key’ and only conflicting keys that contain the word “foo” will be overwritten, any other conflicts would lead to the raise of a `DocumentSyncConflict` exception. A key-strategy may also be a regular expression, so the synchronization above could also be achieved with:

```
dst_job.sync(src_job, doc_sync=sync.DocSync.ByKey('foo'))
```

class `signac.sync.FileSync`

Bases: `object`

Collection of file synchronization strategies.

class `Ask`

Bases: `object`

Resolve sync conflicts by asking whether a file should be overwritten interactively.

static `always(src, dst, fn)`

Resolve sync conflicts by always overwriting.

classmethod `keys()`

Return keys.

static `never(src, dst, fn)`

Resolve sync conflicts by never overwriting.

static `update(src, dst, fn)`

Resolve sync conflicts based on newest modified timestamp.

class `signac.sync.DocSync`

Bases: `object`

Collection of document synchronization functions.

class `ByKey(key_strategy=None)`

Bases: `object`

Synchronize documents key by key.

COPY = `'copy'`

Copy (and potentially overwrite) documents like any other file.

NO_SYNC = `False`

Do not synchronize documents.

static `update(src, dst)`

Perform a simple update.

`signac.sync.sync_jobs(src, dst, strategy=None, exclude=None, doc_sync=None, recursive=False, follow_symlinks=True, preserve_permissions=False, preserve_times=False, preserve_owner=False, preserve_group=False, deep=False, dry_run=False)`

Synchronize the dst job with the src job.

By default, this method will synchronize all files and document data of dst job with the src job until a synchronization conflict occurs. There are two different kinds of synchronization conflicts:

1. The two jobs have files with the same name, but different content.
2. The two jobs have documents that share keys, but those keys are mapped to different values.

A file conflict can be resolved by providing a `FileSync` *strategy* or by *excluding* files from the synchronization. An unresolvable conflict is indicated with the raise of a `FileSyncConflict` exception.

A document synchronization conflict can be resolved by providing a `doc_sync` function that takes the source and the destination document as first and second argument.

Parameters

- **src** (*Job*) – The src job, data will be copied from this job’s workspace.

- **dst** (*Job*) – The dst job, data will be copied to this job’s workspace.
- **strategy** (*callable*) – A synchronization strategy for file conflicts. The strategy should be a callable with signature `strategy(src, dst, filepath)` where `src` and `dst` are the source and destination instances of *Project* and `filepath` is the filepath relative to the project root. If no strategy is provided, a *errors.SyncConflict* exception will be raised upon conflict. (Default value = None)
- **exclude** (*str*) – A filename exclusion pattern. All files matching this pattern will be excluded from the synchronization process. (Default value = None)
- **doc_sync** (attribute or callable from *DocSync*) – A synchronization strategy for document keys. The default is to use a safe key-by-key strategy that will not overwrite any values on conflict, but instead raises a *DocumentSyncConflict* exception.
- **recursive** (*bool*) – Recursively synchronize sub-directories encountered within the job workspace directories. (Default value = False)
- **follow_symlinks** (*bool*) – Follow and copy the target of symbolic links. (Default value = True)
- **preserve_permissions** (*bool*) – Preserve file permissions (Default value = False)
- **preserve_times** (*bool*) – Preserve file modification times (Default value = False)
- **preserve_owner** (*bool*) – Preserve file owner (Default value = False)
- **preserve_group** (*bool*) – Preserve file group ownership (Default value = False)
- **dry_run** (*bool*) – If True, do not actually perform any synchronization operations. (Default value = False)
- **deep** (*bool*) – (Default value = False)

`signac.sync.sync_projects` (*source, destination, strategy=None, exclude=None, doc_sync=None, selection=None, check_schema=True, recursive=False, follow_symlinks=True, preserve_permissions=False, preserve_times=False, preserve_owner=False, preserve_group=False, deep=False, dry_run=False, parallel=False, collect_stats=False*)

Synchronize the destination project with the source project.

Try to clone all jobs from the source to the destination. If the destination job already exist, try to synchronize the job using the optionally specified strategy.

Parameters

- **source** (class:~.*Project*) – The project presenting the source for synchronization.
- **destination** (class:~.*Project*) – The project that is modified for synchronization.
- **strategy** (*callable*) – A synchronization strategy for file conflicts. The strategy should be a callable with signature `strategy(src, dst, filepath)` where `src` and `dst` are the source and destination instances of *Project* and `filepath` is the filepath relative to the project root. If no strategy is provided, a *errors.SyncConflict* exception will be raised upon conflict. (Default value = None)
- **exclude** (*str*) – A filename exclusion pattern. All files matching this pattern will be excluded from the synchronization process. (Default value = None)
- **doc_sync** (attribute or callable from *DocSync*) – A synchronization strategy for document keys. The default is to use a safe key-by-key strategy that will not overwrite any values on conflict, but instead raises a *DocumentSyncConflict* exception.

- **selection** (sequence of *Job* or job ids (str)) – Only synchronize the given selection of jobs. (Default value = None)
- **check_schema** (*bool*) – If True, only synchronize if this and the other project have a matching state point schema. See also: *detect_schema()*. (Default value = True)
- **recursive** (*bool*) – Recursively synchronize sub-directories encountered within the job workspace directories. (Default value = False)
- **follow_symlinks** (*bool*) – Follow and copy the target of symbolic links. (Default value = True)
- **preserve_permissions** (*bool*) – Preserve file permissions (Default value = False)
- **preserve_times** (*bool*) – Preserve file modification times (Default value = False)
- **preserve_owner** (*bool*) – Preserve file owner (Default value = False)
- **preserve_group** (*bool*) – Preserve file group ownership (Default value = False)
- **dry_run** (*bool*) – If True, do not actually perform the synchronization operation, just log what would happen theoretically. Useful to test synchronization strategies without the risk of data loss. (Default value = False)
- **deep** (*bool*) – (Default value = False)
- **parallel** (*bool*) – (Default value = False)
- **collect_stats** (*bool*) – (Default value = False)

Returns Returns stats if `collect_stats` is True, else None.

Return type NoneType or `FileTransferStats`

Raises

- *DocumentSyncConflict* – If there are conflicting keys within the project or job documents that cannot be resolved with the given strategy or if there is no strategy provided.
- *FileSyncConflict* – If there are differing files that cannot be resolved with the given strategy or if no strategy is provided.
- *SchemaSyncConflict* – In case that the `check_schema` argument is True and the detected state point schema of this and the other project differ.

signac.warnings module

Module for signac deprecation warnings.

exception `signac.warnings.SignacDeprecationWarning`

Bases: `UserWarning`

Indicates the deprecation of a signac feature, API or behavior.

This class indicates a user-relevant deprecation and is therefore a `UserWarning`, not a `DeprecationWarning` which is hidden by default.

signac.errors module

Errors raised by signac.

exception `signac.errors.AuthenticationError`
Bases: `signac.core.errors.Error`, `RuntimeError`
Authentication error.

exception `signac.errors.BufferException`
Bases: `signac.core.errors.Error`
An exception occurred in buffered mode.

exception `signac.errors.BufferedFileError` (*files*)
Bases: `signac.core.jsondict.BufferException`
Raised when an error occurred while flushing one or more buffered files.

files
A dictionary of files that caused issues during the flush operation, mapped to a possible reason for the issue or `None` in case that it cannot be determined.

exception `signac.errors.ConfigError`
Bases: `signac.core.errors.Error`, `RuntimeError`
Error with parsing or reading a configuration file.

exception `signac.errors.DestinationExistsError` (*destination*)
Bases: `signac.core.errors.Error`, `RuntimeError`
The destination for a move or copy operation already exists.

Parameters `destination` (*str*) – The destination causing the error.

exception `signac.errors.DocumentSyncConflict` (*keys*)
Bases: `signac.errors.SyncConflict`
Raised when a synchronization operation fails due to a document conflict.

keys = None
The keys that caused the conflict.

exception `signac.errors.Error`
Bases: `Exception`
Base class used for signac Errors.

exception `signac.errors.ExportError`
Bases: `signac.core.errors.Error`, `RuntimeError`
Error exporting documents to a mirror.

exception `signac.errors.FetchError`
Bases: `FileNotFoundError`
Error in fetching data.

exception `signac.errors.FileSyncConflict` (*filename*)
Bases: `signac.errors.SyncConflict`
Raised when a synchronization operation fails due to a file conflict.

filename = None
The filename of the file that caused the conflict.

exception `signac.errors.IncompatibleSchemaVersion`
Bases: `signac.core.errors.Error`
The project's schema version is incompatible with this version of signac.

exception `signac.errors.InvalidKeyError`

Bases: `ValueError`

Raised when a user uses a non-conforming key.

exception `signac.errors.JobsCorruptedError` (*job_ids*)

Bases: `signac.core.errors.Error`, `RuntimeError`

The state point manifest file of one or more jobs cannot be opened or is corrupted.

Parameters `job_ids` – The job id(s) of the corrupted job(s).

exception `signac.errors.KeyTypeError`

Bases: `TypeError`

Raised when a user uses a key of invalid type.

exception `signac.errors.SchemaSyncConflict` (*schema_src*, *schema_dst*)

Bases: `signac.errors.SyncConflict`

Raised when a synchronization operation fails due to schema differences.

exception `signac.errors.StatepointParsingError`

Bases: `signac.core.errors.Error`, `RuntimeError`

Indicates an error that occurred while trying to identify a state point.

exception `signac.errors.SyncConflict`

Bases: `signac.core.errors.Error`, `RuntimeError`

Raised when a synchronization operation fails.

exception `signac.errors.WorkspaceError` (*error*)

Bases: `signac.core.errors.Error`, `OSError`

Raised when there is an issue creating or accessing the workspace.

Parameters `error` – The underlying error causing this issue.

1.3.10 synced_collections package

Data Types

`synced_collections.synced_collection` module

Implement the `SyncedCollection` class.

class `signac.synced_collections.data_types.synced_collection.SyncedCollection` (*parent=None*,
**args*,
***kwargs*)

Bases: `collections.abc.Collection`

An abstract `Collection` type that is synced with a backend.

This class extends `collections.abc.Collection` and adds a number of abstract internal methods that must be implemented by its subclasses. These methods can be split into two groups of functions that are designed to be implemented by separate subtrees in the inheritance hierarchy that can then be composed:

Concrete Collection Types

These subclasses should implement the APIs for specific types of collections. For instance, a list-like `SyncedCollection` should implement the standard methods for sequences. In addition, they must implement the following abstract methods defined by the `SyncedCollection`:

- `is_base_type()`: Determines whether an object satisfies the semantics of the collection object a given *SyncedCollection* is designed to mimic.
- `_to_base()`: Converts a *SyncedCollection* to its natural base type (e.g. a *list*).
- `_update()`: Updates the *SyncedCollection* to match the contents of the provided `collections.abc.Collection`. After calling `sc._update(c)`, we must have that `sc == c`; however, since such updates are frequent when loading and saving data to a resource, `_update()` should be implemented to minimize new object creation wherever possible.

Backend

These subclasses encode the process by which in-memory data is converted into a representation suitable for a particular backend. For instance, a JSON backend should know how to save a Python object into a JSON-encoded file and then read that object back.

- `_load_from_resource()`: Loads data from the underlying resource and returns it in an object satisfying `is_base_type()`.
- `_save_to_resource()`: Stores data to the underlying resource.
- `_backend`: A unique string identifier for the resource backend.

Since these functionalities are effectively completely orthogonal, members of a given group should be interchangeable. For instance, a dict-like *SyncedCollection* can be combined equally easily with JSON, MongoDB, or SQL backends.

Validation

Due to the restrictions of a particular backend or the needs of a particular application, synced collections may need to restrict the data that they can store. Validators provide a standardized mechanism for this. A validator is a callable that parses any data added to a *SyncedCollection* and raises an *Exception* if any invalid data is provided. Validators cannot modify the data and should have no side effects. They are purely provided as a mechanism to reject invalid data. For example, a JSON validator would raise *Exceptions* if it detected non-string keys in a dict.

Since *SyncedCollection* is designed for extensive usage of inheritance, validators may be inherited by subclasses. There are two attributes that subclasses of *SyncedCollection* can define to control the validators used:

- `_validators`: A list of callables that will be inherited by all subclasses.
- `_all_validators`: A list of callables that will be used to validate this class, and this class alone.

When a *SyncedCollection* subclass is initialized (note that this is at *class* definition time, not when instances are created), its `_register_validators()` method will be called. If this class defines an `_all_validators` attribute, this set of validators will be used by all instances of this class. Otherwise, `_register_validators()` will traverse the MRO and collect the `_validators` attributes from all parents of a class, and store these in the `_all_validators` attribute for the class.

Note: Typically, a synced collection will be initialized with resource information, and data will be pulled from that resource. However, initializing with both data and resource information is a valid use case. In this case, the initial data will be validated by the standard validators, however, it will not be checked against the contents stored in the synced resource and is assumed to be consistent. This constructor pattern can be useful to avoid unnecessary resource accesses.

Thread safety

Whether or not *SyncedCollection* objects are thread-safe depends on the implementation of the backend. Thread-safety of *SyncedCollection* objects is predicated on backends providing an atomic write oper-

ation. All concrete collection types use mutexes to guard against concurrent write operations, while allowing read operations to happen freely. The validity of this mode of access depends on the write operations of a `SyncedCollection` being atomic, specifically the `:meth:'~._save_to_resource` method. Whether or not a particular subclass of `SyncedCollection` is thread-safe should be indicated by that subclass setting the `_supports_threading` class variable to `True`. This variable is set to `False` by `SyncedCollection`, so subclasses must explicitly opt-in to support threading by setting this variable to `True`.

Backends that support multithreaded execution will have multithreading support turned on by default. This support can be enabled or disabled using the `enable_multithreading()` and `disable_multithreading()` methods. `enable_multithreading()` will raise a `ValueError` if called on a class that does not support multithreading.

Parameters `parent` (`SyncedCollection`, *optional*) – If provided, the collection within which this collection is nested (Default value = `None`). A parent instance of `SyncedCollection` or `None`. If `None`, the collection owns its own data, otherwise it is nested within its parent. Every `SyncedCollection` either owns its own data, or has a parent (Default value = `None`).

`_LoadSaveType`

alias of `_LoadAndSave`

`_backend`

The backend associated with a given collection.

This property is abstract to enforce that subclasses implement it. Since it's only internal, subclasses can safely override it with just a raw attribute; this property just serves as a way to enforce the abstract API for subclasses.

Type `str`

classmethod `_from_base` (*data*, ***kwargs*)

Dynamically resolve the type of object to the corresponding synced collection.

This method assumes that `data` has already been validated. This assumption can always be met, since this method should only be called internally by other methods that modify the internal collection data. While this requirement does require that all calling methods be responsible for validation, it confers significant performance benefits because it can instruct any invoked class constructors not to validate, which is especially important for nested collections.

Parameters

- **data** (`Collection`) – Data to be converted from base type.
- ****kwargs** – Any keyword arguments to pass to the collection constructor.

Returns Synced object of corresponding base type.

Return type `Collection`

Notes

This method relies on the internal registry of subclasses populated by `__init_subclass__()` and the `is_base_type()` method to determine the subclass with the appropriate backend and data type. Once an appropriate type is determined, that class's constructor is called. Since this method relies on the constructor and other methods, it can be concretely implemented here rather than requiring subclass implementations.

`_load()`

Load the data from the backend.

This method encodes the recursive logic required to handle the loading of nested collections. For a collection contained within another collection, only the root is ever responsible for loading the data. This method handles the appropriate recursive calls, then farms out the actual reading to the abstract method `__load_from_resource()`.

__load_from_resource()

Load data from underlying backend.

This method must be implemented for each backend. Backends may choose to return `None`, signaling that no modification should be performed on the data in memory. This mode is useful for backends where the underlying resource (e.g. a file) may not initially exist, but can be transparently created on save.

Returns An equivalent unsynced collection satisfying `is_base_type()` that contains the data in the underlying resource (e.g. a file).

Return type *Collection* or *None*

classmethod __register_validators()

Register all inherited validators to this class.

This method is called by `__init_subclass__` when subclasses are created to control what validators will be applied to data added to instances of that class. By default, the `__all_validators` class variable defined on the class itself determines the validation rules for that class. If that variable is not defined, then all parents of the class are searched, and a list of validators is constructed by concatenating the `__validators` class variable for each parent class that defines it.

__save()

Save the data to the backend.

This method encodes the recursive logic required to handle the saving of nested collections. For a collection contained within another collection, only the parent is ever responsible for storing the data. This method handles the appropriate recursive calls, then farms out the actual writing to the abstract method `__save_to_resource()`.

__save_to_resource()

Save data to the backend.

This method must be implemented for each backend.

__to_base()

Dynamically resolve the synced collection to the corresponding base type.

This method should not load the data from the underlying resource, it should simply convert the current in-memory representation of a synced collection to its naturally corresponding unsynced collection type.

Returns An equivalent unsynced collection satisfying `is_base_type()`.

Return type *Collection*

__update(data)

Update the in-memory representation to match the provided data.

The purpose of this method is to update the `SyncedCollection` to match the data in the underlying resource. The result of calling this method should be that `self == data`. The reason that this method is necessary is that `SyncedCollections` can be nested, and nested collections must also be instances of `SyncedCollection` so that synchronization occurs even when nested structures are modified. Recreating the full nested structure every time data is reloaded from file is highly inefficient, so this method performs an in-place update that only changes entries that need to be changed.

Parameters `data` (*Collection*) – A collection satisfying `is_base_type()`.

__validate(data)

Validate the input data.

Parameters `data` (`Collection`) – An collection satisfying `is_base_type()`.

classmethod `disable_multithreading()`

Disable all safety checks and thread locks required for thread safety.

The mutex locks required to enable multithreading introduce nontrivial performance costs, so they can be disabled for classes that support it.

classmethod `enable_multithreading()`

Enable safety checks and thread locks required for thread safety.

Support for multithreaded execution can be disabled by calling `disable_multithreading()`; calling this method reverses that.

classmethod `is_base_type(data)`

Check whether data is of the same base type (such as list or dict) as this class.

Parameters `data` (`Any`) – The input data to test.

Returns Whether or not the object can be converted into this synced collection type.

Return type `bool`

class `signac.synced_collections.data_types.synced_collection._LoadAndSave(collection)`

Bases: `object`

A context manager for `SyncedCollection` to wrap saving and loading.

Any write operation on a synced collection must be preceded by a load and followed by a save. Moreover, additional logic may be required to handle other aspects of the synchronization, particularly the acquisition of thread locks. This class abstracts this concept, making it easy for subclasses to customize the behavior if needed (for instance, to introduce additional locks).

synced_collections.synced_dict module

Implements the `SyncedDict`.

This implements a dict-like data structure that also conforms to the `SyncedCollection` API and can be combined with any backend type to give a dict-like API to a synchronized data structure.

class `signac.synced_collections.data_types.synced_dict.SyncedDict(data=None, _validate=True, *args, **kwargs)`

Bases: `signac.synced_collections.data_types.synced_collection.SyncedCollection`, `collections.abc.MutableMapping`

Implement the dict data structure along with values access through attributes named as keys.

The `SyncedDict` inherits from `SyncedCollection` and `MutableMapping`. Therefore, it behaves like a dict.

Parameters

- **data** (`Mapping`, *optional*) – The initial data to populate the dict. If `None`, defaults to `{}` (Default value = `None`).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the `SyncedDict` object behaves like a `dict`, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying backend, copying (even deep copying) a `SyncedDict` instance may exhibit unexpected behavior. If a true copy is required, you should use the `_to_base()` method to get a `dict` representation, and if necessary construct a new `SyncedDict`.

clear () → None. Remove all items from D.

get (*k*, *d*) → D[*k*] if *k* in D, else *d*. *d* defaults to None.

classmethod is_base_type (*data*)

Check whether the data is an instance of mapping.

Parameters *data* (*any*) – Data to be checked.

Returns

Return type `bool`

items () → a set-like object providing a view on D’s items

keys () → a set-like object providing a view on D’s keys

pop (*k*, *d*) → *v*, remove specified key and return the corresponding value.
If key is not found, *d* is returned if given, otherwise `KeyError` is raised.

popitem () → (*k*, *v*), remove and return some (key, value) pair
as a 2-tuple; but raise `KeyError` if D is empty.

reset (*data*)

Update the instance with new data.

Parameters *data* (*mapping*) – Data to update the instance.

Raises `ValueError` – If the data is not a mapping.

setdefault (*k*, *d*) → D.get(*k*,*d*), also set D[*k*]=*d* if *k* not in D

update (*[E]*, ***F*) → None. Update D from mapping/iterable E and F.

If E present and has a `.keys()` method, does: for *k* in E: D[*k*] = E[*k*] If E present and lacks `.keys()` method, does: for (*k*, *v*) in E: D[*k*] = *v* In either case, this is followed by: for *k*, *v* in F.items(): D[*k*] = *v*

values () → an object providing a view on D’s values

synced_collections.synced_list module

Implements the `SyncedList`.

This implements a list-like data structure that also conforms to the `SyncedCollection` API and can be combined with any backend type to give a list-like API to a synchronized data structure.

```
class signac.synced_collections.data_types.synced_list.SyncedList (data=None,
                                                                _validate=True,
                                                                *args,
                                                                **kwargs)
```

Bases: `signac.synced_collections.data_types.synced_collection.SyncedCollection`, `collections.abc.MutableSequence`

Implementation of list data structure.

The `SyncedList` inherits from `SyncedCollection` and `MutableSequence`. Therefore, it behaves similar to a `list`.

Parameters

- **data** (*Sequence, optional*) – The initial data to populate the list. If `None`, defaults to `[]` (Default value = `None`).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the `SyncedList` object behaves like a `list`, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying backend, copying (even deep copying) a `SyncedList` instance may exhibit unexpected behavior. If a true copy is required, you should use the `_to_base()` method to get a `list` representation, and if necessary construct a new `SyncedList`.

`append` (*item*)

`S.append(value)` – append value to the end of the sequence

`clear` () → `None` – remove all items from `S`

`extend` (*iterable*)

`S.extend(iterable)` – extend sequence by appending elements from the iterable

`insert` (*index, item*)

`S.insert(index, value)` – insert value before index

`classmethod is_base_type` (*data*)

Check whether the data is an non-string `Sequence`.

Parameters **data** (*Any*) – Data to be checked

Returns

Return type `bool`

`remove` (*value*)

`S.remove(value)` – remove first occurrence of value. Raise `ValueError` if the value is not present.

`reset` (*data*)

Update the instance with new data.

Parameters **data** (*non-string Sequence*) – Data to update the instance.

Raises `ValueError` – If the data is not a non-string sequence.

Backends

`synced_collections.backends.collection_json` module

Implements a JSON `SyncedCollection` backend.

```
class signac.synced_collections.backends.collection_json.BufferedJSONAttrDict (filename=None,
                                                                    write_concern=Fa
                                                                    data=None,
                                                                    par-
                                                                    ent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: `signac.synced_collections.backends.collection_json.BufferedJSONDict`,
`signac.synced_collections.data_types.attr_dict.AttrDict`

A buffered `JSONAttrDict`.

class `signac.synced_collections.backends.collection_json.BufferedJSONAttrList` (`filename=None`,
`write_concern=False`,
`data=None`,
`parent=None`,
`*args`,
`**kwargs`)

Bases: `signac.synced_collections.backends.collection_json.BufferedJSONList`

A `BufferedJSONList` whose dict-like children will be of type `BufferedJSONAttrDict`.

class `signac.synced_collections.backends.collection_json.BufferedJSONCollection` (`parent=None`,
`filename=None`,
`*args`,
`**kwargs`)

Bases: `signac.synced_collections.backends.serialized_file_buffered_collection.SerializedFileBufferedCollection`,
`signac.synced_collections.backends.collection_json.JSONCollection`

A `JSONCollection` that supports I/O buffering.

This class implements the buffer protocol defined by `BufferedCollection`. The concrete implementation of buffering behavior is defined by the `SerializedFileBufferedCollection`.

class `signac.synced_collections.backends.collection_json.BufferedJSONDict` (`filename=None`,
`write_concern=False`,
`data=None`,
`parent=None`,
`*args`,
`**kwargs`)

Bases: `signac.synced_collections.backends.collection_json.BufferedJSONCollection`,
`signac.synced_collections.data_types.synced_dict.SyncedDict`

A buffered `JSONDict`.

class `signac.synced_collections.backends.collection_json.BufferedJSONList` (`filename=None`,
`write_concern=False`,
`data=None`,
`parent=None`,
`*args`,
`**kwargs`)

Bases: `signac.synced_collections.backends.collection_json.BufferedJSONCollection`,
`signac.synced_collections.data_types.synced_list.SyncedList`

A buffered `JSONList`.

```
class signac.synced_collections.backends.collection_json.JSONAttrDict (filename=None,
                                                                    write_concern=False,
                                                                    data=None,
                                                                    parent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: `signac.synced_collections.backends.collection_json.JSONDict`, `signac.synced_collections.data_types.attr_dict.AttrDict`

A dict-like data structure that synchronizes with a persistent JSON file.

Unlike `JSONAttrDict`, this class also supports attribute-based access to dictionary contents, e.g. `doc.foo == doc['foo']`.

Examples

```
>>> doc = JSONAttrDict('data.json', write_concern=True)
>>> doc['foo'] = "bar"
>>> assert doc.foo == doc['foo'] == "bar"
>>> assert 'foo' in doc
>>> del doc['foo']
>>> doc['foo'] = dict(bar=True)
>>> doc
{'foo': {'bar': True}}
>>> doc.foo.bar = False
>>> doc
{'foo': {'bar': False}}
```

Parameters

- **filename** (*str*, *optional*) – The filename of the associated JSON file on disk (Default value = None).
- **write_concern** (*bool*, *optional*) – Ensure file consistency by writing changes back to a temporary file first, before replacing the original file (Default value = False).
- **data** (`collections.abc.Mapping`, *optional*) – The initial data passed to `JSONAttrDict`. If None, defaults to {} (Default value = None).
- **parent** (`JSONCollection`, *optional*) – A parent instance of `JSONCollection` or None. If None, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the `JSONAttrDict` object behaves like a `dict`, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying file, copying (even deep copying) a `JSONAttrDict` instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new `JSONAttrDict` instance.

```
class signac.synced_collections.backends.collection_json.JSONAttrList (filename=None,
                                                                    write_concern=False,
                                                                    data=None,
                                                                    par-
                                                                    ent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: *signac.synced_collections.backends.collection_json.JSONList*

A *JSONList* whose dict-like children will be of type *JSONAttrDict*.

```
class signac.synced_collections.backends.collection_json.JSONCollection (filename=None,
                                                                    write_concern=False,
                                                                    *args,
                                                                    **kwargs)
```

Bases: *signac.synced_collections.data_types.synced_collection.SyncedCollection*

SyncedCollection

A *SyncedCollection* that synchronizes with a JSON file.

This collection implements synchronization by reading and writing the associated JSON file in its entirety for every read/write operation. This backend is a good choice for maximum accessibility and transparency since all data is immediately accessible in the form of a text file with no additional tooling, but is likely a poor choice for high performance applications.

Thread safety

The *JSONCollection* is thread-safe. To make these collections safe, the `write_concern` flag is ignored in multithreaded execution, and the write is **always** performed via a write to temporary file followed by a replacement of the original file. The file replacement operation uses `os.replace()`, which is guaranteed to be atomic by the Python standard.

Parameters

- **filename** (*str*) – The filename of the associated JSON file on disk.
- **write_concern** (*bool, optional*) – Ensure file consistency by writing changes back to a temporary file first, before replacing the original file (Default value = False).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

filename

The name of the associated JSON file on disk.

Type *str*

```
class signac.synced_collections.backends.collection_json.JSONDict (filename=None,
                                                                    write_concern=False,
                                                                    data=None,
                                                                    par-
                                                                    ent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: *signac.synced_collections.backends.collection_json.JSONCollection*, *signac.synced_collections.data_types.synced_dict.SyncedDict*

A dict-like data structure that synchronizes with a persistent JSON file.

Examples

```
>>> doc = JSONDict('data.json', write_concern=True)
>>> doc['foo'] = "bar"
>>> assert doc['foo'] == "bar"
>>> assert 'foo' in doc
>>> del doc['foo']
>>> doc['foo'] = dict(bar=True)
>>> doc
{'foo': {'bar': True}}
```

Parameters

- **filename** (*str*, *optional*) – The filename of the associated JSON file on disk (Default value = None).
- **write_concern** (*bool*, *optional*) – Ensure file consistency by writing changes back to a temporary file first, before replacing the original file (Default value = False).
- **data** (*collections.abc.Mapping*, *optional*) – The initial data passed to *JSONDict*. If None, defaults to {} (Default value = None).
- **parent** (*JSONCollection*, *optional*) – A parent instance of *JSONCollection* or None. If None, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the *JSONDict* object behaves like a *dict*, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying file, copying (even deep copying) a *JSONDict* instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new *JSONDict* instance.

```
class signac.synced_collections.backends.collection_json.JSONList (filename=None,
                                                                    write_concern=False,
                                                                    data=None,
                                                                    par-
                                                                    ent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: *signac.synced_collections.backends.collection_json.JSONCollection*,
signac.synced_collections.data_types.synced_list.SyncedList

A list-like data structure that synchronizes with a persistent JSON file.

Only non-string sequences are supported by this class.

Examples

```
>>> synced_list = JSONList('data.json', write_concern=True)
>>> synced_list.append("bar")
>>> assert synced_list[0] == "bar"
>>> assert len(synced_list) == 1
>>> del synced_list[0]
```

Parameters

- **filename** (*str*, *optional*) – The filename of the associated JSON file on disk (Default value = None).
- **write_concern** (*bool*, *optional*) – Ensure file consistency by writing changes back to a temporary file first, before replacing the original file (Default value = None).
- **data** (non-str `collections.abc.Sequence`, *optional*) – The initial data passed to `JSONList``. If `None``, defaults to `[]`` (Default value = None).
- **parent** (`JSONCollection`, *optional*) – A parent instance of `JSONCollection`` or `None``. If `None``, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the `JSONList`` object behaves like a `list``, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying file, copying (even deep copying) a `JSONList`` instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new `JSONList`` instance.

```
class signac.synced_collections.backends.collection_json.MemoryBufferedJSONAttrDict (filename=
write_con
data=None
par-
ent=None
*args,
**kwargs
```

```
Bases:
    signac.synced_collections.backends.collection_json.
    MemoryBufferedJSONDict,    signac.synced_collections.data_types.attr_dict.
    AttrDict
```

A buffered `JSONAttrDict``.

```
class signac.synced_collections.backends.collection_json.MemoryBufferedJSONAttrList (filename=
write_con
data=None
par-
ent=None
*args,
**kwargs
```

```
Bases:
    signac.synced_collections.backends.collection_json.
    MemoryBufferedJSONList
```

A `MemoryBufferedJSONList`` whose dict-like children will be of type `MemoryBufferedJSONAttrDict``.

```
class signac.synced_collections.backends.collection_json.MemoryBufferedJSONCollection (parent
file-
name=
*args,
**kwa
```

```
Bases:
    signac.synced_collections.buffered.memory_buffered_collection.
    SharedMemoryFileBufferedCollection,    signac.synced_collections.backends.
    collection_json.JSONCollection
```

A *JSONCollection* that supports I/O buffering.

This class implements the buffer protocol defined by *BufferedCollection*. The concrete implementation of buffering behavior is defined by the *SharedMemoryFileBufferedCollection*.

```
class signac.synced_collections.backends.collection_json.MemoryBufferedJSONDict (filename=None,
                                                                              write_concern=
                                                                              data=None,
                                                                              par-
                                                                              ent=None,
                                                                              *args,
                                                                              **kwargs)
```

```
Bases:
    signac.synced_collections.backends.collection_json.
    MemoryBufferedJSONCollection,
    signac.synced_collections.data_types.
    synced_dict.SyncedDict
```

A buffered *JSONDict*.

```
class signac.synced_collections.backends.collection_json.MemoryBufferedJSONList (filename=None,
                                                                              write_concern=
                                                                              data=None,
                                                                              par-
                                                                              ent=None,
                                                                              *args,
                                                                              **kwargs)
```

```
Bases:
    signac.synced_collections.backends.collection_json.
    MemoryBufferedJSONCollection,
    signac.synced_collections.data_types.
    synced_list.SyncedList
```

A buffered *JSONList*.

```
signac.synced_collections.backends.collection_json.json_attr_dict_validator (data)
Validate data for JSONAttrDict.
```

This validator combines the logic from the following validators into one to make validation more efficient:

This validator combines the following logic:

- JSON format validation
- Ensuring no dots are present in string keys
- Converting non-str keys to strings. This is a backwards compatibility layer that will be removed in signac 2.0.

Parameters *data* – Data to validate.

Raises

- *KeyTypeError* – If key data type is not supported.
- *TypeError* – If the data type of *data* is not supported.

synced_collections.backends.collection_mongodb module

Implements a MongoDB *SynchedCollection* backend.

```
class signac.synced_collections.backends.collection_mongodb.MongoDBCollection (collection=None,
                                                                    uid=None,
                                                                    parent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: `signac.synced_collections.data_types.synced_collection.SyncedCollection`

A `SyncedCollection` that synchronizes with a MongoDB document.

In MongoDB, a database is composed of multiple MongoDB **collections**, which are analogous to tables in SQL databases but do not enforce a schema like in relational databases. In turn, collections are composed of **documents**, which are analogous to rows in a table but are much more flexible, storing any valid JSON object in a JSON-like encoded format known as BSON (“binary JSON”).

Each `MongoDBCollection` can be represented as a MongoDB document, so this backend stores the `MongoDBCollection` as a single document within the collection provided by the user. The document is identified by a unique key provided by the user.

Thread safety

The `MongoDBCollection` is not thread-safe.

Parameters

- **collection** (`pymongo.collection.Collection`) – The MongoDB client in which to store data.
- **uid** (`dict`) – The unique key-value mapping added to the data and stored in the document so that it is uniquely identifiable in the MongoDB collection. The key “data” is reserved and may not be part of this uid.
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: The user is responsible for providing a unique id such that there are no possible collisions between different `MongoDBCollection` instances stored in the same MongoDB collection. Failure to do so may result in data corruption if multiple documents are found to be apparently associated with a given uid.

collection

Get the collection being synced to.

Type `pymongo.collection.Collection`

uid

Get the unique mapping used to identify this collection.

Type `dict`

```
class signac.synced_collections.backends.collection_mongodb.MongoDBDict (collection=None,
                                                                    uid=None,
                                                                    data=None,
                                                                    parent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: `signac.synced_collections.backends.collection_mongodb.`

MongoDBCollection, *signac.synced_collections.data_types.synced_dict*,
SyncedDict

A dict-like data structure that synchronizes with a document in a MongoDB collection.

Examples

```
>>> doc = MongoDBDict('data')
>>> doc['foo'] = "bar"
>>> assert doc['foo'] == "bar"
>>> assert 'foo' in doc
>>> del doc['foo']
>>> doc['foo'] = dict(bar=True)
>>> doc
{'foo': {'bar': True}}
```

Parameters

- **collection** (*pymongo.collection.Collection*, optional) – A *pymongo.collection.Collection* instance (Default value = None).
- **uid** (*dict*, optional) – The unique key-value mapping identifying the collection (Default value = None).
- **data** (non-str *collections.abc.Mapping*, optional) – The initial data passed to *MongoDBDict*. If None, defaults to {} (Default value = None).
- **parent** (*MongoDBCollection*, optional) – A parent instance of *MongoDBCollection* or None. If None, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the *MongoDBDict* object behaves like a *dict*, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying database, copying a *MongoDBDict* instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new *MongoDBDict* instance.

```
class signac.synced_collections.backends.collection_mongodb.MongoDBList (collection=None,
                                                                    uid=None,
                                                                    data=None,
                                                                    par-
                                                                    ent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: *signac.synced_collections.backends.collection_mongodb*.
MongoDBCollection, *signac.synced_collections.data_types.synced_list*.
SyncedList

A list-like data structure that synchronizes with a document in a MongoDB collection.

Only non-string sequences are supported by this class.

Examples

```
>>> synced_list = MongoDBList('data')
>>> synced_list.append("bar")
>>> assert synced_list[0] == "bar"
>>> assert len(synced_list) == 1
>>> del synced_list[0]
```

Parameters

- **collection** (*pymongo.collection.Collection*, optional) – A *pymongo.collection.Collection* instance (Default value = None).
- **uid** (*dict*, optional) – The unique key-value mapping identifying the collection (Default value = None).
- **data** (non-str *collections.abc.Sequence*, optional) – The initial data passed to *MongoDBList*. If None, defaults to [] (Default value = None).
- **parent** (*MongoDBCollection*, optional) – A parent instance of *MongoDBCollection* or None. If None, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the *MongoDBList* object behaves like a *list*, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying database, copying a *MongoDBList* instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new *MongoDBList* instance.

synced_collections.backends.collection_redis module

Implements a Redis *SyncedCollection* backend.

```
class signac.synced_collections.backends.collection_redis.RedisCollection (client=None,  
key=None,  
*args,  
**kwargs)
```

Bases: *signac.synced_collections.data_types.synced_collection.SyncedCollection*

A *SyncedCollection* that synchronizes with a Redis database.

This backend stores data in Redis by associating it with the provided key.

Thread safety

The *RedisCollection* is not thread-safe.

Parameters

- **client** (*redis.Redis*) – The Redis client used to persist data.
- **key** (*str*) – The key associated with this collection in the Redis database.
- ***args** – Positional arguments forwarded to parent constructors.

- ****kwargs** – Keyword arguments forwarded to parent constructors.

client

The Redis client used to store the data.

Type *redis.Redis*

key

The key associated with this collection stored in Redis.

Type *str*

```
class signac.synced_collections.backends.collection_redis.RedisDict (client=None,
                                                                    key=None,
                                                                    data=None,
                                                                    par-
                                                                    ent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: *signac.synced_collections.backends.collection_redis.RedisCollection*,
signac.synced_collections.data_types.synced_dict.SyncedDict

A dict-like data structure that synchronizes with a persistent Redis database.

Examples

```
>>> doc = RedisDict('data')
>>> doc['foo'] = "bar"
>>> assert doc['foo'] == "bar"
>>> assert 'foo' in doc
>>> del doc['foo']
>>> doc['foo'] = dict(bar=True)
>>> doc
{'foo': {'bar': True}}
```

Parameters

- **client** (*redis.Redis*, *optional*) – A redis client (Default value = None).
- **key** (*str*, *optional*) – The key of the collection (Default value = None).
- **data** (*collections.abc.Mapping*, *optional*) – The initial data passed to *RedisDict*. If None, defaults to {} (Default value = None).
- **parent** (*RedisCollection*, *optional*) – A parent instance of *RedisCollection* or None. If None, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the *RedisDict* object behaves like a *dict*, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying database, copying a *RedisDict* instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new *RedisDict* instance.

```
class signac.synced_collections.backends.collection_redis.RedisList (client=None,
                                                                    key=None,
                                                                    data=None,
                                                                    par-
                                                                    ent=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: `signac.synced_collections.backends.collection_redis.RedisCollection`,
`signac.synced_collections.data_types.synced_list.SyncedList`

A list-like data structure that synchronizes with a persistent Redis database.

Only non-string sequences are supported by this class.

Examples

```
>>> synced_list = RedisList('data')
>>> synced_list.append("bar")
>>> assert synced_list[0] == "bar"
>>> assert len(synced_list) == 1
>>> del synced_list[0]
```

Parameters

- **client** (`redis.Redis`, *optional*) – A Redis client (Default value = None).
- **key** (`str`, *optional*) – The key of the collection (Default value = None).
- **data** (non-str `collections.abc.Sequence`, *optional*) – The initial data passed to `RedisList`. If None, defaults to [] (Default value = None).
- **parent** (`RedisCollection`, *optional*) – A parent instance of `RedisCollection` or None. If None, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the `RedisList` object behaves like a `list`, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying database, copying a `RedisList` instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new `RedisList` instance.

synced_collections.backends.collection_zarr module

Implements a Zarr `SyncedCollection` backend.

```
class signac.synced_collections.backends.collection_zarr.ZarrCollection (group=None,
                                                                    name=None,
                                                                    codec=None,
                                                                    *args,
                                                                    **kwargs)
```

Bases: `signac.synced_collections.data_types.synced_collection.SyncedCollection`

A *SyncedCollection* that synchronizes with a Zarr group.

Since Zarr is designed for storage of array-like data, this backend implements synchronization by storing the collection in a 1-element object array. The user provides the group within which to store the data and the name of the data in the group.

Thread safety

The *ZarrCollection* is not thread-safe.

Parameters

- **group** (*zarr.hierarchy.Group*) – The Zarr group in which to store data.
- **name** (*str*) – The name under which this collection is stored in the Zarr group.
- **codec** (*numcodecs.abc.Codec*) – The encoding mechanism for the data. If not provided, defaults to JSON encoding (Default value: None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

codec

The encoding method used for the data.

Type *numcodecs.abc.Codec*

group

The Zarr group storing the data.

Type *zarr.hierarchy.Group*

name

The name of this data in the Zarr group.

Type *str*

```
class signac.synced_collections.backends.collection_zarr.ZarrDict (group=None,
                                                                name=None,
                                                                data=None,
                                                                parent=None,
                                                                *args,
                                                                **kwargs)
```

Bases: *signac.synced_collections.backends.collection_zarr.ZarrCollection*,
signac.synced_collections.data_types.synced_dict.SyncedDict

A dict-like data structure that synchronizes with a Zarr group.

Examples

```
>>> doc = ZarrDict('data')
>>> doc['foo'] = "bar"
>>> assert doc['foo'] == "bar"
>>> assert 'foo' in doc
>>> del doc['foo']
>>> doc['foo'] = dict(bar=True)
>>> doc
{'foo': {'bar': True}}
```

Parameters

- **group** (*zarr.hierarchy.Group*, optional) – The group in which to store data (Default value = None).
- **name** (*str*, optional) – The name of the collection (Default value = None).
- **data** (*collections.abc.Mapping*, optional) – The initial data passed to *ZarrDict*. If None, defaults to {} (Default value = None).
- **parent** (*ZarrCollection*, optional) – A parent instance of *ZarrCollection* or None. If None, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the *ZarrDict* object behaves like a `dict`, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying database, copying (even deep copying) a *ZarrDict* instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new *ZarrDict* instance.

```
class signac.synced_collections.backends.collection_zarr.ZarrList (group=None,
                                                             name=None,
                                                             data=None,
                                                             par-
                                                             ent=None,
                                                             *args,
                                                             **kwargs)
```

Bases: *signac.synced_collections.backends.collection_zarr.ZarrCollection*,
signac.synced_collections.data_types.synced_list.SyncedList

A list-like data structure that synchronizes with a Zarr group.

Only non-string sequences are supported by this class.

Examples

```
>>> synced_list = ZarrList('data')
>>> synced_list.append("bar")
>>> assert synced_list[0] == "bar"
>>> assert len(synced_list) == 1
>>> del synced_list[0]
```

Parameters

- **group** (*zarr.hierarchy.Group*, optional) – The group in which to store data (Default value = None).
- **name** (*str*, optional) – The name of the collection (Default value = None).
- **data** (non-str *collections.abc.Sequence*, optional) – The initial data passed to *ZarrList*. If None, defaults to [] (Default value = None).
- **parent** (*ZarrCollection*, optional) – A parent instance of *ZarrCollection* or None. If None, the collection owns its own data (Default value = None).
- ***args** – Positional arguments forwarded to parent constructors.
- ****kwargs** – Keyword arguments forwarded to parent constructors.

Warning: While the `ZarrList` object behaves like a `list`, there are important distinctions to remember. In particular, because operations are reflected as changes to an underlying database, copying (even deep copying) a `ZarrList` instance may exhibit unexpected behavior. If a true copy is required, you should use the call operator to get a dictionary representation, and if necessary construct a new `ZarrList` instance.

Buffers

`synced_collections.buffers.buffered_collection` module

Defines a buffering protocol for `SyncedCollection` objects.

Depending on the choice of backend, synchronization may be an expensive process. In that case, it can be helpful to allow many in-memory modifications to occur before any synchronization is attempted. Since many collections could be pointing to the same underlying resource, maintaining proper data coherency across different instances requires careful consideration of how the data is stored. The appropriate buffering methods can differ for different backends; as a result, the basic interface simply lays out the API for buffering and leaves implementation details for specific backends to handle. Judicious use of buffering can dramatically speed up code paths that might otherwise involve, for instance, heavy I/O. The specific buffering mechanism must be implemented by each backend since it depends on the nature of the underlying data format.

All buffered collections expose a local context manager for buffering. In addition, each backend exposes a context manager `BufferedCollection.buffer_backend()` that indicates to all buffered collections of that backend that they should enter buffered mode. These context managers may be nested freely, and buffer flushes will occur when all such managers have been exited.

class `signac.synced_collections.buffers.buffered_collection.BufferedCollection` (*args, **kwargs)
 Bases: `signac.synced_collections.data_types.synced_collection.SyncedCollection`

A `SyncedCollection` defining an interface for buffering.

The default behavior of this class is not to buffer. This class simply defines an appropriate interface for buffering behavior so that client code can rely on these methods existing, e.g. to be able to do things like `with collection.buffered...`. This feature allows client code to indicate to the collection when it is safe to buffer reads and writes, which usually means guaranteeing that the synchronization destination (e.g. an underlying file or database entry) will not be modified by other processes concurrently with the set of operations within the buffered block. However, in the default case the result of this will be a no-op and all data will be immediately synchronized with the backend.

The `BufferedCollection` overrides the `_load()` and `_save()` methods to check whether buffering is enabled or not. If not, the behavior is identical to the parent class. When in buffered mode, however, the `BufferedCollection` introduces two additional hooks that can be overridden by subclasses to control how the collection behaves while buffered:

- `_load_from_buffer()`: Loads data while in buffered mode and returns it in an object satisfying `is_base_type()`. The default behavior is to simply call `_load_from_resource()`.
- `_save_to_buffer()`: Stores data while in buffered mode. The default behavior is to simply call `_save_to_resource()`.

Thread safety

Whether or not buffering is thread safe depends on the buffering method used. In general, both the buffering logic and the data type operations must be thread safe for the resulting collection type to be thread safe.

`__flush()`

Flush data associated with this instance from the buffer.

`classmethod __flush_buffer()`

Flush all data in this class's buffer.

`__is_buffered`

Check if we should write to the buffer or not.

`__load()`

Load data from the backend but buffer if needed.

This method is identical to the *SyncedCollection* implementation except that it determines whether data is actually synchronized or instead read from a temporary buffer based on the buffering mode.

`__load_from_buffer()`

Store data in buffer.

By default, this method simply calls *__load_from_resource()*. Subclasses must implement specific buffering strategies.

Returns An equivalent unsynced collection satisfying *is_base_type()* that contains the buffered data. By default, the buffered data is just the data in the resource.

Return type *Collection*

`__save()`

Synchronize data with the backend but buffer if needed.

This method is identical to the *SyncedCollection* implementation for *sync* except that it determines whether data is actually synchronized or instead written to a temporary buffer based on the buffering mode.

`__save_to_buffer()`

Store data in buffer.

By default, this method simply calls *__save_to_resource()*. Subclasses must implement specific buffering strategies.

`classmethod backend_is_buffered()`

Check if this backend is currently buffered.

`classmethod buffer_backend(*args, **kwargs)`

Enter context to buffer all operations for this backend.

synced_collections.buffers.file_buffered_collection module

A standardized buffering implementation for file-based backends.

All file-based backends can use a similar buffering protocol. In particular, integrity checks can be performed by checking for whether the file has been modified since it was originally loaded into the buffer. However, various specific components are abstract and must be implemented by child classes.

class `signac.synced_collections.buffers.file_buffered_collection.FileBufferedCollection` (*partially implemented*)
file-based buffered collection
name: file_buffered_collection
args: file, name, mode
kwargs: kwargs

Bases: `signac.synced_collections.buffers.buffered_collection.BufferedCollection`

A *SyncedCollection* that can buffer file I/O.

This class provides a standardized buffering protocol for all file-based backends. All file-based backends can use the same set of integrity checks prior to a buffer flush to ensure that no conflicting modifications are made. Specifically, they can check whether the file has been modified on disk since it was originally loaded to the buffer. This class provides the basic infrastructure for that and defines standard methods that can be used by all classes. Subclasses must define the appropriate storage mechanism.

Note:

Important notes for developers:

- This class should be inherited before any other collections. This requirement is due to the extensive use of multiple inheritance. Since this class is designed to be combined with other *SyncedCollection* types without making those types aware of buffering behavior, it transparently hooks into the initialization process, but this is dependent on its constructor being called before those of other classes.
 - All subclasses must define a class level `_BUFFER_CAPACITY` variable that is used to determine the maximum allowable buffer size.
-

Parameters `filename` (*str*, *optional*) – The filename of the associated JSON file on disk (Default value = None).

Warning: Although it can be done safely, in general modifying two different collections pointing to the same underlying resource while both are in different buffering modes is unsupported and can lead to undefined behavior. This class makes a best effort at performing safe modifications, but it is possible to construct nested buffered contexts for different objects that can lead to an invalid buffer state, or even situations where there is no obvious indicator of what is the canonical source of truth. In general, if you need multiple objects pointing to the same resource, it is **strongly** recommended to work with both of them in identical buffering states at all times.

classmethod `buffer_backend` (*buffer_size=None*, *force_write=None*, **args*, ***kwargs*)

Enter context to buffer all operations for this backend.

Parameters

- **buffer_size** (*int*) – The capacity of the buffer to use within this context (resets after the context is exited).
- **force_write** (*bool*) – This argument does nothing and is only present for compatibility with signac 1.x.

classmethod `disable_multithreading` ()

Disable all safety checks and thread locks required for thread safety.

This method adds managed buffer-related thread safety in addition to what the parent method does.

classmethod `enable_multithreading` ()

Enable safety checks and thread locks required for thread safety.

This method adds managed buffer-related thread safety in addition to what the parent method does.

classmethod `get_buffer_capacity` ()

Get the current buffer capacity.

Returns The amount of data that can be stored before a flush is triggered in the appropriate units for a particular buffering implementation.

Return type `int`

classmethod `get_current_buffer_size()`

Get the total amount of data currently stored in the buffer.

Returns The size of all data contained in the buffer in the appropriate units for a particular buffering implementation.

Return type `int`

classmethod `set_buffer_capacity(new_capacity)`

Update the buffer capacity.

Parameters `new_capacity(int)` – The new capacity of the buffer in the appropriate units for a particular buffering implementation.

`synced_collections.buffers.serialized_file_buffered_collection` module

Buffering for file-based backends using a serialized buffer.

The buffering method implemented here involves a single buffer of serialized data. All collections in buffered mode encode their data into this buffer on save and decode from it on load.

class `signac.synced_collections.buffers.serialized_file_buffered_collection.SerializedFile`

Bases: `signac.synced_collections.buffers.file_buffered_collection.FileBufferedCollection`

A `FileBufferedCollection` based on a serialized data store.

This class extends the `FileBufferedCollection` and implements a concrete storage mechanism in which data is encoded (by default, into JSON) and stored into a buffer. This buffer functions as a central data store for all collections and is a synchronization point for various collections pointing to the same underlying file. This serialization method may be a bottleneck in some applications; see the Warnings section for more information.

The buffer size and capacity for this class is measured in the total number of bytes stored in the buffer that correspond to file data. This is *not* the total size of the buffer, which also contains additional information like the hash of the data and the file metadata (which are used for integrity checks), but it is the relevant metric for users.

Note: Important note for subclasses: This class should be inherited before any other collections. This requirement is due to the extensive use of multiple inheritance: since this class is designed to be combined with other `SyncedCollection` types without making those types aware of buffering behavior, it transparently hooks into the initialization process, but this is dependent on its constructor being called before those of other classes.

Thread safety

This buffering method is thread safe. This thread safety is independent of the safety of an individual collection backend; the backend must support thread safe writes to the underlying resource in order for a buffered version using this class to be thread safe for general use. The thread safety guaranteed by this class only concerns buffer reads, writes, and flushes. All these operations are serialized because there is no way to prevent one collection from triggering a flush while another still thinks its data is in the cache.

Parameters `filename(str, optional)` – The filename of the associated JSON file on disk (Default value = None).

Warning:

- Although it can be done safely, in general modifying two different collections pointing to the same underlying resource while both are in different buffering modes is unsupported and can lead to undefined behavior. This class makes a best effort at performing safe modifications, but it is possible to construct nested buffered contexts for different objects that can lead to an invalid buffer state, or even situations where there is no obvious indicator of what is the canonical source of truth. In general, if you need multiple objects pointing to the same resource, it is **strongly** recommended to work with both of them in identical buffering states at all times.
- The overhead of this buffering method is quite high due to the constant encoding and decoding of data. For performance-critical applications where memory is not highly constrained and virtual memory limits are absent, the *SharedMemoryFileBufferedCollection* may be more appropriate.
- Due to the possibility of read operations triggering a flush, the contents of the buffer may be invalidated on loads as well. To prevent this even nominally read-only operations are serialized. As a result, although this class is thread safe, it will effectively serialize all operations and will therefore not be performant.

synced_collections.buffered.memory_buffered_collection module

A standardized buffering implementation for file-based backends.

The buffering method implemented here involves a single buffer of references to in-memory objects containing data. These objects are the base types of a given *SyncedCollection* type, e.g. a dict for all dict-like collections, and are the underlying data stores for those types. This buffering method exploits the fact that all mutable collection types in Python are references, so modifying one such collection results in modifying all of them, thereby removing any need for more complicated synchronization protocols.

class `signac.synced_collections.buffered.memory_buffered_collection.SharedMemoryFileBufferedCollection`

Bases: `signac.synced_collections.buffered.file_buffered_collection.FileBufferedCollection`

A *SyncedCollection* that defers all I/O when buffered.

This class extends the *FileBufferedCollection* and implements a concrete storage mechanism in which collections store a reference to their data in a buffer. This method takes advantage of the reference-based semantics of built-in Python mutable data types like dicts and lists. All collections referencing the same file are pointed to the same underlying data store in buffered mode, allowing all changes in one to be transparently reflected in the others. To further improve performance, the buffer size is determined only based on the number of modified collections stored, not the total number. As a result, the maximum capacity is only reached when a large number of modified collections are stored, and unmodified collections are only removed from the buffer when a buffered context is exited (rather than when buffer capacity is exhausted). See the Warnings section for more information.

The buffer size and capacity for this class is measured in the total number of collections stored in the buffer that have undergone any modifications since their initial load from disk. A sequence of read-only operations will load data into the buffer, but the apparent buffer size will be zero.

Note: Important note for subclasses: This class should be inherited before any other collections. This requirement is due to the extensive use of multiple inheritance: since this class is designed to be combined with other

SyncedCollection types without making those types aware of buffering behavior, it transparently hooks into the initialization process, but this is dependent on its constructor being called before those of other classes.

Thread safety

This buffering method is thread safe. This thread safety is independent of the safety of an individual collection backend; the backend must support thread safe writes to the underlying resource in order for a buffered version using this class to be thread safe for general use. The thread safety guaranteed by this class only concerns buffer reads, writes, and flushes. All these operations are serialized because there is no way to prevent one collection from triggering a flush while another still thinks its data is in the cache; however, this shouldn't be terribly performance-limiting since in buffered mode we're avoiding I/O anyway and that's the only thing that can be effectively parallelized here.

Parameters `filename` (*str*, *optional*) – The filename of the associated JSON file on disk (Default value = None).

Warning:

- Although it can be done safely, in general modifying two different collections pointing to the same underlying resource while both are in different buffering modes is unsupported and can lead to undefined behavior. This class makes a best effort at performing safe modifications, but it is possible to construct nested buffered contexts for different objects that can lead to an invalid buffer state, or even situations where there is no obvious indicator of what is the canonical source of truth. In general, if you need multiple objects pointing to the same resource, it is **strongly** recommended to work with both of them in identical buffering states at all times.
- This buffering method has no upper bound on the buffer size if all operations on buffered objects are read-only operations. If a strict upper bound is required, for instance due to strict virtual memory limits on a given system, use of the *SerializedFileBufferedCollection* will allow limiting the total memory usage of the process.

Miscellaneous Modules

`synced_collections.utils` module

Define common utilities.

```
class signac.synced_collections.utils.AbstractTypeResolver (abstract_type_identifiers,  
                                                         cache_blocklist=None)
```

Bases: `object`

Mapping between recognized types and their abstract parents.

Synced collections are heavily reliant on checking the types of objects to determine the appropriate type of behavior in various scenarios. For maximum generality, most of these checks use the ABCs defined in `collections.abc`. The price of this flexibility is that *isinstance* checks with these classes are very slow because the `__instancecheck__` hooks are implemented in pure Python and require checking many different cases.

Rather than attempting to directly optimize this behavior, this class provides a workaround by which we can amortize the cost of type checks. Given a set of types that must be resolved and a way to identify each of these (which may be expensive), it maintains a local cache of all instances of a given type that have previously been observed. This reduces the cost of type checking to a simple `dict` lookup, except for the first time a new type is observed.

Parameters

- **abstract_type_identifiers** (*Mapping*) – A mapping from a string identifier for a group of types (e.g. "MAPPING") to a callable that can be used to identify that type. Due to insertion order guarantees of dictionaries in Python>=3.6 (officially 3.7), it may be beneficial to order this dictionary with the most frequently occurring types first. However, unless users have many different concrete types implementing the same abstract interface (e.g. many Mapping types identified via `isinstance(obj, Mapping)`), any performance gain should be negligible since the callables will only be executed once per type.
- **cache_blocklist** (*Sequence, optional*) – A sequence of string identifiers from `abstract_type_identifiers` that should not be cached. If there are cases where objects of the same type would be classified into separate groups based on the callables in `abstract_type_identifiers`, this argument allows users to specify that this type should not be cached. This argument should be used sparingly because performance will quickly degrade if many calls to `get_type()` are with types that cannot be cached. The identifiers (keys in `abstract_type_identifiers`) corresponding to elements of the blocklist should be placed first in the `abstract_type_identifiers` dictionary since they will never be cached and are therefore the most likely callables to be used repeatedly (Default value = None).

abstract_type_identifiers

A mapping from string identifiers for an abstract type to callables that accepts an object and returns True if the object is of the key type and False if not.

Type Dict[str, Callable[Any, bool]]

type_map

A mapping from concrete types to the corresponding named abstract type from `abstract_type_identifiers`.

Type Dict[Type, str]

get_type (*obj*)

Get the type string corresponding to this data type.

Parameters **obj** (*Any*) – Any object whose type to check

Returns The name of the type, where valid types are the keys of the dict argument to the constructor. If the object's type cannot be identified, will return None.

Return type str

```
class signac.synced_collections.utils.SyncedCollectionJSONEncoder(*, skip-
    keys=False, en-
    sure_ascii=True, check_circular=True, al-
    low_nan=True, sort_keys=False, in-
    dent=None, separa-
    tors=None, de-
    fault=None)
```

Bases: `json.encoder.JSONEncoder`

A `json.JSONEncoder` that handles objects encodeable using `default()`.

Warning:

- JSON encoding of numpy arrays is not invertible; once encoded, reloading the data will result in converting arrays to lists and numpy numbers into ints or floats.
- This class assumes that the in-memory data for a SyncedCollection is up-to-date. If the data has been changed on disk without updating the collection, or if this class is used to serialize the data before any method of the collection is invoked that would load the data from disk, the resulting serialized data may be incorrect.

default (*o: Any*) → Dict[str, Any]

Implement this method in a subclass such that it returns a serializable object for *o*, or calls the base implementation (to raise a `TypeError`).

For example, to support arbitrary iterators, you could implement `default` like this:

```
def default(self, o):
    try:
        iterable = iter(o)
    except TypeError:
        pass
    else:
        return list(iterable)
    # Let the base class default method raise the TypeError
    return JSONEncoder.default(self, o)
```

`signac.synced_collections.utils.default` (*o: Any*) → Dict[str, Any]

Get a JSON-serializable version of compatible types.

This function is suitable for use with JSON-serialization tools as a way to serialize `SyncedCollection` objects and NumPy arrays. It will attempt to obtain a JSON-serializable representation of an object that is otherwise not serializable by attempting to access its `_data` attribute.

Warning:

- JSON encoding of numpy arrays is not invertible; once encoded, reloading the data will result in converting arrays to lists and numpy numbers into ints or floats.
- This function assumes that the in-memory data for a SyncedCollection is up-to-date. If the data has been changed on disk without updating the collection, or if this function is used to serialize the data before any method is invoked that would load the data from disk, the resulting serialized data may be incorrect.

synced_collections.validators module

Validators for SyncedCollection API.

A validator is any callable that raises Exceptions when called with invalid data. Validators should act recursively for nested data structures and should not return any values, only raise errors. This module implements built-in validators, but client code is free to implement and add additional validators to collection types as needed.

`signac.synced_collections.validators.json_format_validator` (*data*)

Validate input data can be serialized to JSON.

Parameters `data` – Data to validate.

Raises

- `KeyTypeError` – If key data type is not supported.
- `TypeError` – If the data type of data is not supported.

`signac.synced_collections.validators.no_dot_in_key(data)`

Raise an exception if there is a dot (.) in a mapping's key.

Parameters `data` – Data to validate.

Raises

- `KeyTypeError` – If key data type is not supported.
- `InvalidKeyError` – If the key contains invalid characters or is otherwise malformed.

`signac.synced_collections.validators.require_string_key(data)`

Raise an exception if key in a mapping is not a string.

Almost all supported backends require string keys.

Parameters `data` – Data to validate.

Raises `KeyTypeError` – If key type is not a string.

1.4 Changelog

The **signac** package follows semantic versioning.

1.4.1 Version 1

[1.7.0] – 2021-06-08

Added

- New `SyncedCollection` class and subclasses to replace `JSONDict` with more general support for different types of resources (such as MongoDB collections or Redis databases) and more complete support for different data types synchronized with files (#196, #234, #249, #316, #383, #397, #465, #484, #529, #530). This change introduces a minor-backwards incompatible change; for users making direct use of signac buffering, the `force_write` parameter is no longer respected. If the argument is passed, a warning will now be raised to indicate that it is ignored and will be removed in signac 2.0.
- Unified querying for state point and document filters using 'sp' and 'doc' as prefixes (#332, #514). This change introduces a minor backwards-incompatible change to the `Collection` index schema ('statepoint'->'sp'), but this does not affect any APIs, only indexes saved to file using a previous version of signac. Indexing APIs will be removed in signac 2.0.

Changed

- Optimized internal path joins to speed up project iteration (#515).

Deprecated

- `doc_filter` arguments, which are replaced by namespaced filters. Due to their long history, `doc_filter` arguments will still be accepted in signac 2.0 and will only be removed in 3.0 (#516).
- The modules `signac.core.attrdict`, `signac.core.json`, `signac.core.jsondict`, and `signac.core.syncdict.py` are deprecated in favor of the new `SyncedCollection` classes and will be removed in signac 2.0 (#483).

Fixed

- Corrected docstrings for `Job.update_statepoint` and `Project.update_statepoint` (#506, #563).

[1.6.0] – 2021-01-24

Added

- Implemented `JobsCursor.__contains__` check (#449).
- Added documentation for `JobsCursor` class (#475).

Changed

- Optimized job hash and equality checks (#442, #455).
- Optimized `H5Store` initialization (#443).
- State points are loaded lazily when `Job` is opened by id (#238, #239).
- Optimized `Job` and `Project` classes to cache internal properties and initialize on access (#451).
- Python 3.6 is only tested with oldest dependencies (#474).
- Improved documentation for updating and resetting state points (#444).

Deprecated

- Deprecate `syncutil.copytree` method (#439).

Fixed

- Zero-dimensional NumPy arrays can be used in state points and documents (#449).

[1.5.1] – 2020-12-19

Added

- Support for h5py version 3 (#411).
- Added `pyupgrade` to pre-commit hooks (#413).

- Code is formatted with `black` and `isort` pre-commit hooks (#415).
- Added macOS to CircleCI testing pipeline (#281, #414).
- Official support for Python 3.9 (#417).

Changed

- Optimized internal function `_mkdir_p` (#421).
- Optimized performance of job initialization (#422).
- Optimized performance of buffer storage (#428).
- Optimized performance of creating/loading synced data structures (#429).

[1.5.0] – 2020-09-20

Added

- Type annotations are validated during continuous integration (#313).
- Added `__repr_html__` method in `ProjectSchema` class (#314, #324).
- Allow grouping by variables that are not present in all jobs in the project in `JobsCursor.groupby` (#321, #323).
- Added parameters `usecols` and `flatten` to allow selection of columns and flattening of nested data when converting signac data into a pandas `DataFrame` (#327, #330).
- Added support for `pre-commit` hooks (#355, #358).
- Expanded CLI documentation (#187, #359, #377).

Changed

- Docstrings are now written in `numpydoc` style.

Fixed

- Fix the `signac config verify` command (previously broken) (#301, #302).
- Warnings now appear when raised by the `signac` CLI (#317, #308).
- Fix dots in synchronization error messages (#375, #376).

Deprecated

- Deprecate the `create_access_modules` method in `Project`, to be removed in 2.0 (#303, #308).
- The `MainCrawler` class has replaced the `MasterCrawler` class. Both classes are deprecated (#342).

Removed

- Dropped support for Python 3.5 (#340). The signac project will follow the [NEP 29 deprecation policy](#) going forward.
- Removed dependency on `pytest-subtests` (#379).

[1.4.0] – 2020-02-28

Added

- Added Windows to platforms tested with continuous integration (#264, #266).
- Add command line option `-m/--merge` for `signac sync` (#280, #230).

Changed

- Workspace directory is created when `Project` is initialized (#267, #271).
- Changed testing framework from `unittest` to `pytest` (#212, #275).
- Refactored internal use of deprecated `get_statepoint` function (#227, #282).

Fixed

- Fixed issues on Windows with `H5Store`, project import/export, and operations that move files (#264, #266).
- Calling `items` or `values` on `_SyncedDict` objects does not mutate nested dictionaries (#234, #269).
- Fixed issue with `project.data` access from separate instances of `H5StoreManager` (#274, #278).
- Fixed error when launching `signac shell` if permissions are denied for `.signac_shell_history` (#279).

Removed

- Removed vendored `tqdm` module and replaced it with a requirement (#289).
- Removed support for `rapidjson` as an alternative JSON library (#285, #287).
- Removed tuple of keys implementation of nested dictionaries (#272, #296).

[1.3.0] – 2019-12-20

Added

- Official support for Python 3.8 (#258).
- Add properties `Project.id` and `Job.id` (#250).
- Add `signac.diff_jobs` function and `$ signac diff` CLI command to compare two or more state points (#248, #247).
- Add function to initialize a sample data space for testing purposes (#215).

- Add schema version to ensure compatibility and enable migrations in future package versions (#165, #253).

Changed

- Implemented `Project.__contains__` check in constant time (#231).

Fixed

- Attempting to create a linked view for a Project on Windows now raises an informative error message (#214, #236).
- Project configuration is initialized using `ConfigObj`, allowing the configuration to include commas and special characters (#251, #252).

Deprecated

- Deprecate the `get_id` method in `Project` and `Job` classes in favor of the `id` property, to be removed in 2.0 (#250).
- In-memory modification of the project configuration, to be removed in 2.0 (#246).

Removed

- Dropped support for Python 2.7 (#232).

[1.2.0] – 2019-07-22

Added

- Keep signac shell command history on a per-project basis (#134, #194).
- Add `read_json()` and `to_json()` methods to `Collection` class (#104, #200).

Fixed

- Fix issue where shallow copies of instances of `Job` would behave incorrectly (#153, #207).
- Fix issue causing a failure of the automatic conversion of valid key types (#168, #205).
- Improve the ‘dots in keys’ error message to make it easier to fix related issues (#170, #205).
- Update the `__repr__` and `__repr_html__` implementations of the `Project`, `Job`, and `JobsCursor` classes (#193).
- Reduce the logging verbosity about a missing default host key in the configuration (#201).
- Fix issue with incorrect detection of dict-like files managed with the `DictManager` class (e.g. `job.stores`) (#203).
- Fix issue with generating views from the command line for projects with only one job (#208, #211).
- Fix issue with heterogeneous types in state point values that are lists (#209, #210).

Deprecated

- Support for Python 2.7 is deprecated with this version and will be removed in the next minor release. See the [Python 3 Statement](#).

[1.1.0] – 2019-05-19

Added

- Add command line options `--sp` and `--doc` for `signac find` that allow users to display key-value pairs of the state point and document in combination with the job id (#97, #146).
- Improve the representation (return value of `repr()`) of instances of `H5Group` and `SyncedAttrDict`.

Fixed

- Fix: Searches for whole numbers will match all numerically matching integers regardless of whether they are stored as decimals or whole numbers (#169).
- Fix: Passing an instance of `dict` to `H5Store.setdefault()` will return an instance of `H5Group` instead of a `dict` (#180).
- Fix error with storing numpy arrays and scalars in a synced dictionary (e.g. `job.statepoint`, `job.document`) (#184).
- Fix issue with `ResourceWarning` originating from unclosed instance of `Collection` (#186).
- Fix issue with using the `get_project()` function with a relative path and `search=False` (#191).

Removed

- Support for Python version 3.4 (no longer tested).

[1.0.0] – 2019-02-28

Added

- Official support for Python 3.7.
- The `H5Store` and `H5StoreManager` classes, which are useful for storing (numerical) array-like data with an HDF5-backend. These classes are exposed within the root namespace.
- The `job.data` and `project.data` properties which present an instance of `H5Store` to access numerical data within the job workspace and project root directory.
- The `job.stores` and `project.stores` properties, which present an instance of `H5StoreManager` to manage multiple instances of `H5Store` to store numerical array-like data within the project workspace and project root directory.
- The `signac.get_job()` and the `signac.Project.get_job()` functions that allow users to get a job handle by switching into or providing the job's workspace directory.
- The `job` variable is automatically set when opening a `signac shell` from within a job's workspace directory.

- Add the `signac shell -c` option which allows the direct specification of Python commands to be executed within the shell.
- Automatic cast of `numpy` arrays to lists when storing them within a `JSONDict`, e.g., a `job.statepoint` or `job.document`.
- Enable `Collection` class to manage collections stored in compressed files (`gzip`, `zip`, etc.).
- Enable deleting of `JSONDict` keys through the attribute interface, e.g., `del job.doc.foo`.
- Pretty HTML representation of instances of `Project` and `JobsCursor` targeted at Jupyter Notebooks (requires `pandas`, automatically enabled when installed).
- The `to_dataframe()` function to export the job state point and document data of a `Project` or a `JobsCursor`, e.g., the result of `Project.find_jobs()`, as a `pandas.DataFrame` (requires `pandas`).

Changed

- Dots (`.`) in keys are no longer allowed for `JSONDict` and `Collection` keys (previously deprecated).
- The `JSONDict` module is exposed in the root namespace, which is useful for storing text-serializable data with a JSON-backend similar to the `job.statepoint` or `job.document`, etc.
- The `Job.init()` method returns the job to allow one-line job creation and initialization.
- The `search` argument was added to the `signac.get_project()` function, which when `True` (the default), will cause `signac` to search for a project within *and above* a specified root directory, not only within the root directory. The behavior without any arguments remains unchanged.

Fixed

- Fix `Collection.update()` behavior such that existing documents with identical primary key are updated. Previously, a `KeyError` would be raised.
- Fix issue where the `Job.move()` would trigger a confusing `DestinationExists` exception when trying to move jobs across devices / file systems.
- Fix issue that caused failures when the `python-rapidjson` package is installed. The `python-rapidjson` package is used as the primary JSON-backend when installed.
- Fix issue where schema with multiple keys would subset incorrectly if the list of jobs or statepoints was provided as an iterator rather than a sequence.

Removed

- Removes the obsolete and deprecated `core.search_engine` module.
- The previously deprecated `Project.find_statepoints()` and `Project.find_job_documents()` functions have been removed.
- The `Project.find_jobs()` no longer accepts the obsolete `index` argument.

1.4.2 Version 0.9

[0.9.5] – 2019-01-31

Fixed

- Ensure that the `next()` function can be called for a `JobsIterator`, e.g., `project.find()`.
- Pickling issue that occurs when a `_SyncedDict` (`job.statepoint`, `job.document`, etc.) contains a list.
- Issue with the `readline` module that would cause `signac shell` to fail on Windows operating systems.

[0.9.4] – 2018-10-24

Added

- Adds the `$ signac import` command and the `Project.import_from()` method for the import of data spaces into a project workspace, such as a directory, a tarball, or a zip file.
- Adds the `$ signac export` command and the `Project.export_to()` method for the export of project workspaces to an external location, such as a directory, a tarball, or a zip file.
- Adds functionality for the rapid initialization of temporary projects with the `signac.TemporaryProject` context manager.
- Adds the `signac.Project.temporary_project()` context manager which creates a temporary project within the root project workspace.
- Add `signac` to the default namespace when invoking `signac shell`.
- Add option to specify a custom view path for the `signac view/Project.create_linked_view()` function.
- Iterables of documents used to construct a `Collection` no longer require an `_id` field.

Changed

- The default path for linked views has been adjusted to match the one used for data exports.

Fixed

- Fix issue where differently typed integer values stored within a `Collection` under the same key would not be indexed correctly. This issue affected the correct function of the `$type` operator for aforementioned cases and would lead to incorrect types in the `Project` schema detection algorithm for integer values.
- Fix issue where jobs that are migrated (state point change), but are not initialized, were not properly updated.
- Fix issue where changes to lists as part of synchronized dictionary, for example a state point or document would not be saved.
- Fix non-deterministic issue occurring on network file systems when trying to open jobs where the user has no write access to the job workspace directory.

[0.9.3] – 2018-06-14

Added

- Add `$near` operator to express queries for numerical values that match up to a certain precision.
- Add the `$ signac shell` sub command to directly launch a Python interpreter within a project directory.

Fixed

- Fix issue where a job instance would not be properly updated after more than one state point reset.

[0.9.2] – 2017-12-18

Added

- Add provisional feature (persistent state point caching); calling the `Project.update_cache()` method will generate and store a persistent state point cache in the project root directory, which will increase the speed of many project iteration, search, and selection operations.
- Add `Project.check()` method which checks for workspace corruption, but does not make any attempt to repair it.
- The `Project.repair()` method will attempt to repair jobs, that have been corrupted by manually renaming the job's workspace directory.

Changed

- Enable the `write_concern` flag for the `job.document`.
- Allow to omit the specification of an authentication mechanism in the MongoDB host configuration.

Fixed

- Fix critical issue in the `JSONDict` implementation that would previously overwrite the underlying file when attempting to store values that are not JSON serializable.
- Fix issue where the `Project.export()` function would ignore the update argument when the index to export to would be a MongoDB collection.

[0.9.1] – 2017-11-07

Fixed

- Fix critical issue in the `SyncedAttrDict` implementation that would previously overwrite the underlying file if the first operation was a `__setitem__()` operation.

[0.9.0] – 2017-10-28

Added

- Introduction of `$ signac sync`, `Project.sync()`, and `Job.sync()` for the simplified and fine-grained synchronization of multiple project data spaces.
- Introduction of `$ signac schema` and `Project.detect_schema()` for the automatic detection of the implicit and semi-structured state point schema of a project data space.
- Simplified aggregation of jobs over projects and `Project.find_jobs()` results with the `Project.groupby()` function.

- Support for project-centralized data with the `Project.document` attribute and the `Project.fn()` method for the wrapping of filenames within the project root directory.
- Added the `Job.clear()` and the `Job.reset()` methods to clear or reset a job's workspace data.

Changed

- Both `Job.statepoint` and `Job.document` now use the same underlying data structure and provide the exact same API (copy with `()` and access of keys as attributes).
- The `Collection` class uses an internal counter instead of UUIDs for the generation of primary keys (resulting in improved performance).
- Major performance improvements (faster `Collection`, improved caching)
- Overhaul of the reference documentation.

1.4.3 Version 0.8

[0.8.7] – 2017-10-05

Fixed

- Fix an issue where the creation of linked views was non-deterministic in some cases.
- Fix an issue where the creation of linked views would fail when the project contains job with state points that have lists as values.

[0.8.6] – 2017-08-25

Fixed

- Fix `Collection` append truncation issue (see issue #66).

[0.8.5] – 2017-06-07

Changed

- The `signac` ids in the `signac find -show` view are no longer enclosed by quotation marks.

Fixed

- Fix compatibility issue that broke the `signac find -view` and all `-pretty` commands on Python 2.7.
- Fix issue where view directories would be incomplete in combination with heterogeneous state point schemas.

[0.8.4] – 2017-05-19**Added**

- All search queries on project and collection objects support various operators including: *\$and*, *\$or*, *\$gt*, *\$gte*, *\$lt*, *\$lte*, *\$eq*, *\$ne*, *\$exists*, *\$regex*, *\$where*, *\$in*, *\$nin*, and *\$type*.
- The `$ signac find` command supports a simple filter syntax, where key value pairs can be provided as individual arguments.
- The `$ signac find` command is extended by a `-show` option, to display the state point and the document contents directly. The contents are truncated to an adjustable depth to reduce output noise.
- The `$ signac view` command has an additional filter option to select a sub data space directly without needing to pipe job ids.
- The new `$ signac document` command can be used to display a job's document directly.

Changed

- Minor performance improvements.

[0.8.3] – 2017-05-10**Changed**

- Raise `ExportError` when updating with an empty index.

Fixed

- Fix command line logic issue with `$signac config host`.
- Fix bug, where `Collection.replace_one()` would ignore the `upsert` argument under specific conditions.

[0.8.2] – 2017-04-19**Fixed**

- Fixes a `TypeError` which occurred under specific conditions when calling `Collection.find()` with nested filter arguments.

[0.8.1] – 2017-04-17**Fixed**

- Fixes wide-spread typo (*indeces* -> *indexes*).

[0.8.0] – 2017-04-16

Overall major simplification of the generation of indexes and the management and searching of index collections without external database.

Added

- Introduction of the `Collection` class for the management of document collections, such as indexes in memory and on disk.
- Generation of file indexes directly via the `signac.index_files()` function.
- Generation of main indexes directly via the `signac.index()` function and the `$ signac index` command.
- The API of `signac_access.py` files has been simplified, including the possibility to use a blank file for a minimal configuration.
- Use the `$ signac project --access` command to create a minimal access module in addition to `Project.create_access_module()`.
- The update of existing index collections has been simplified by using the `export()` function with the `update=True` argument, which means that stale documents (the associated file or state point no longer exists) are automatically identified and removed.
- Added the `Job.ws` attribute, as short-cut for `Job.workspace()`.
- The `Job.sp` interface has a `get()` function which can be used to specify a default value in case that the requested key is not part of the state point.

Changed (breaking API)

- The `$ signac index` command generates a main index instead of a project index. To generate a project index from the command line use `$ signac project --index` instead.
- The `SignacProjectCrawler` class expects the project's root directory as first argument, not the workspace directory.
- The `get_crawlers()` function defined within a `signac_access.py` access module is expected to yield crawler instances directly, not a mapping of crawler ids and instances.
- The simplification of the `signac_access.py` module API is reflected in a reduction of arguments to the `Project.create_access_module()` method.

Changed (non-breaking)

- The `RegexFileCrawler`, `SignacProjectCrawler` and `MainCrawler` classes were moved into the root namespace.
- If a `MainCrawler` object is instantiated with the `raise_on_error` argument set to `True`, any errors encountered during crawling are raised instead of ignored and skipped; this simplifies the debugging of erroneous access modules.
- Improved error message for invalid configuration files.
- Better error messages for invalid `$ signac find` queries.
- Check a host configuration on the command line via `$ signac host --test`.
- A MongoDB database host configuration defaults to `none` when no authentication method is explicitly specified.
- Using the `--debug` option in combination with `$ signac index` will show the traceback of errors encountered during indexing instead of ignoring them.
- Instances of `Job` are hashable, making it possible to use them as dict keys for instance.

- The representation of `Job` instances via `repr()` can actually serves as copy constructor command.
- The project interface implementation performs all non-trivial search operations on an internally management index collection, which improves performance and simplifies the code base.

Deprecated

- The `DocumentSearchEngine` class has been deprecated, its functionality is now provided by the `Collection` class.

Fixed

- An issue related to exporting documents to MongoDB collections via `pymongo` in combination with Python 2.7 has been fixed.

1.4.4 Version 0.7

[0.7.1] – 2017-01-09

Added

- When the `python-rapidjson` package is installed, it will be used for JSON encoding/decoding (experimental).

Changed

- All job move-related methods raise `DestinationExistsError` in case of destination conflicts.
- Optimized `$ signac find` command.

Fixed

- Fixed bug in `$ signac statepoint`.
- Suppress 'broken pipe error' message when using `$ signac find` for example in combination with `$ head`.

[0.7.0] – 2017-01-04

Added

- Add support for Python version 3.6.
- Add support for PyPy and PyPy3.
- Simplified iteration over project data spaces.
- An existing linked view can be updated by executing the view command again.
- Add attribute interface for the access and modification of job state points: `Job.sp`.
- Add function for moving and copying of jobs between projects.

- All project related iterators support the `len`-operator.
- Enable iteration over all jobs with: `for job in project:`
- Make `len(project)` an alias for `project.num_jobs()`.
- Add `in`-operator to determine whether a job is initialized within a project.
- Add `Job.sp` attribute to access and modify a job's state point.
- The `Project.open_job()` method accepts abbreviated job ids.
- Add `Project.min_len_unique_id()` method to determine the minimum length of job ids to be unique within the project's data space.
- Add `Job.move()` method to move jobs between projects.
- Add `Project.clone()` method to copy jobs between projects.
- Add `$ signac move` and `$ signac clone` command line functions.
- Add `Job.reset_statepoint()` method to reset a job's state point.
- Add `Job.update_statepoint()` method to update a job's state point.
- Add a `Job.FN_DOCUMENT` constant which defines the default filename of the job document file
- The `$ signac find` command accepts a `-d/--doc-filter` option to filter by job document contents.
- Add the `Project.create_linked_view()` method as replacement for the previously deprecated `Project.create_view()` method.

Changed

- Linked views use relative paths.
- The *Guide* documentation chapter has been renamed to *Reference* and generally overhauled.
- The *Quick Reference* documentation chapter has been extended.

Fixed

- Fix error when using an instance of `Job` after calling `Job.remove()`.
- A project created in one the standard config directories (such as the home directory) does not take prevalence over project configurations in or above the current working directory.

Removed

- The *signac-gui* component has been removed.
- The `Project.create_linked_view()` `force` argument is removed.
- The `Project.find_variable_parameters()` method has been removed

1.4.5 Version 0.6

[0.6.2] – 2017-12-15

Added

- Add instructions on how to acknowledge **signac** in publications to documentation.
- Add cite module for the auto-generation of formatted references and BibTeX entries.

Removed

- Remove SSL authentication support.

[0.6.1] – 2017-11-26

Changed

- The `Project.create_view()` method triggers a `DeprecationWarning` instead of a `PendingDeprecationWarning`.
- The `Project.find_variable_parameters()` method triggers a `DeprecationWarning` instead of a `PendingDeprecationWarning`.

Fixed

- Make package more robust against PySide import errors.
- Fix `Project.__repr__` method.
- Fix critical bug in `fs.GridFS` class, which rendered it unusable.
- Fix issue in `indexing.fetch()` previously resulting in local paths being ignored.
- Fix error `signac.__all__` namespace directive.

[0.6.0] – 2016-11-18

Added

- Add the `export_to_mirror()` function for mirroring files.
- Introduction of the `signac.fs` namespace to simplify the configuration of mirror filesystems.
- Add `errors` module to root namespace. Many exceptions raised inherit from the base exception types defined within that module, making it easier to catch signac related errors.
- Add the `export_one()` function for the export of a single index document; simplifies the implementation of custom export functions.
- Opening an instance of `Job` with the `open_job()` method multiple times and entering a job context recursively is now well-defined behavior: Opening a job now adds the current working directory onto a stack, closing it switches into the directory on top of the stack.

- The return type of `Project.open_job()` can be configured to make it easier to specialize projects with custom job types.

Changed

- The `MainCrawler` logic has been simplified; their primary function is the compilation of index documents from subcrawlers, all export logic, including data mirroring is now provided by the `signac.export()` function.
- Each index document is now uniquely coupled with only one file or data object, which is why `signac.fetch()` replaces `signac.fetch_one()` and the latter one has been deprecated and is currently an alias of the former one.
- The `signac.fetch()` function always returns a file-like object, regardless of format definition.
- The format argument in the `crawler.define()` function is now optional and has now very well defined behavior for str types. It is encouraged to define a format with a str constant rather than a file-like object type.
- The `TextFile` file-like object class definition in the `formats` module has been replaced with a constant of type `str`.
- The `signac.export()` function automatically delegates to specialized implementations such as `export_pymongo()` and is more robust against errors, such as broken connections.
- The `export_pymongo()` function makes multiple automatic restart attempts when encountering errors.
- Documentation: The tutorial is now based on `signac-examples` jupyter notebooks.
- The `contrib.crawler` module has been renamed to `contrib.indexing` to better reflect the semantic context.
- The `signac.export()` function now implements the logic for data linking and mirroring.
- Provide default argument for `'-indent'` option for `$ signac statepoint` command.
- Log, but do not reraise exceptions during `MainCrawler` execution, making the compilation of main indexes more robust against errors.
- The object representation of `Job` and `Project` instances is simplified.
- The warning verbosity has been reduced when importing modules with optional dependencies.

Removed

- All modules related to the stale *conversion framework* feature have been removed resulting in a removal of the optional `networkx` dependency.
- Multiple modules related to the *conversion framework* feature have been removed, including: `contrib.formats_network`, `contrib.conversion`, and `contrib.adapters`.

Fixed

- Opening instances of `Job` with the `Job.open()` method multiple times, equivalently entering the job context recursively, does not cause an error anymore, but instead the behavior is well-defined.

1.4.6 Version 0.5

[0.5.0] – 2016-08-31

Added

- New function: `signac.init_project()` simplifies project initialization within Python
- Added optional `root` argument to `signac.get_project()` to simplify getting a project handle outside of the current working directory
- Added optional argument to `signac.get_project()`, to allow fetching of projects outside of the current working directory.
- Added two class factory methods to *Project*: `get_project()` and `init_project()`.

Changed

- The performance of project indexing and crawling has been improved.

1.4.7 Version 0.4

[0.4.0] – 2016-08-05

Added

- The performance of find operations can be greatly improved by using pre-generated job indexes.
- New top-level commands: `$ signac find`, `$ signac index`, `$ signac statepoint`, and `$ signac view`.
- New method: `Project.create_linked_view()`
- New method: `Project.build_job_statepoint_index()`
- New method: `Project.build_job_search_index()`
- The `Project.find_jobs()` method allows to filter by job document.

Changed

- The `SignacProjectCrawler` indexes all jobs, not only those with non-empty job documents.
- The `signac.fetch_one()` function returns `None` if no associated object can be fetched.
- The tutorial is restructured into multiple parts.
- Instructions for installation are separated from the guide.

Removed

- Remove previously deprecated `crawl` keyword argument in index export functions.
- Remove previously deprecated function `common.config.write_config()`.

1.4.8 Version 0.3

[0.3.0] – 2016-06-23

Added

- Add contributing agreement and guidelines.

Changed

- Change license from MIT to BSD 3-clause license.

1.4.9 Version 0.2

[0.2.9] – 2016-06-06

Added

- Addition of the `signac config` command line API.
- Password updates are encrypted with `bcrypt` when `passlib` is installed.
- The user is prompted to enter missing credentials (username/password) in case that they are not stored in the configuration.
- The `$ signac config` tool provides the `--update-pw` argument, which allows users to update their own password.
- Added MIT license, in addition, all source code files contain a short licensing header.

Changed

- Improved documentation on how to configure `signac`.
- The OSI classifiers are updated, including an upgrade of the development status to '4 - beta'.

Fixed

- Nested job state points can no longer get corrupted. This bug occurred when trying to operate on nested state point mappings.

Deprecated

- Deprecated `pymongo` versions 2.x are no longer supported.

[0.2.8] – 2016-04-18**Added**

- `Project` is now in the root namespace.
- Add `index()` method to *Project*.
- Add `create_access_module()` method to *Project*.
- Add `find_variable_parameters()` method to *Project*.
- Add `fn()` method to *Job*, which prepends the job's workspace path to a filename.
- The documentation contains a comprehensive tutorial

Changed

- The `crawl()` function yields only the index documents and not a tuple of `(_id, doc)`.
- `export()` and `export_pymongo()` expect the index documents as first argument, not a crawler instance. The old API is still supported, but will trigger a `DeprecationWarning`.

[0.2.7] – 2016-02-29**Added**

- Add `job.isfile()` method

Changed

- Optimize `project.find_statepoints()` and `project.repair()` functions.

[0.2.6] – 2016-02-20**Added**

- Add `job.reset_statepoint()` and `job.update_statepoint()`
- Add `job.remove()` function

Changed

- Sanitize filter argument in all `project.find_*` methods.
- The `job.statepoint()` function accurately represents saved statepoints, e.g. tuples are represented as lists, as there is no difference between tuples and lists in JSON.
- `signac-gui` does not block on database operations.
- `signac-gui` allows reload of databases and collections of connected hosts.

Fixed

- `RegexFileCrawler.define()` class function only acts upon the actual specialization and not globally on all `RegexFileCrawler` classes.
- `signac-gui` does not crash when replica sets are configured.

[0.2.5] – 2016-02-10

Added

- Added `signac.get_project()`, `signac.get_database()`, `signac.fetch()` and `signac.fetch_one()` to top-level namespace.
- Added basic shell commands, see `$ signac --help`.
- Allow opening of jobs by id: `project.open_job(id='abc123...')`.
- Mirror data while crawling.
- Use extra sources for `fetch()` and `fetch_one()`.
- Add file system handler: `LocalFS`, handler for local file system.
- Add file system handler: `GridFS`, handler for MongoDB GridFS file system.
- Crawler tags, to control which crawlers are used for a specific index.
- Allow explicit job workspace creation with `job.init()`.
- Forwarding of pymongo host configuration via **signac** configuration.

Changed

- Major reorganization of the documentation, split into: Overview, Guide, Quick Reference and API.
- Documentation: Add notes for system administrators about advanced indexing.
- Warn about outdated pymongo versions.
- Set `zip_safe` flag to true in `setup.py`.
- Remove dependency on `six` module, by adding it to the common subpackage.

Deprecated

Fixed

- Fixed hard import of pymongo bug (issue #24).
- Crawler issues with malformed documents.

[0.2.4] – 2016-01-11

Added

- Implement `Project.repair()` function for projects with corrupted workspaces.

- Allow environment variables in workspace path definition.
- Check and fix config permission errors.

Changed

- Increase robustness of job manifest file creation.

Fixed

- Fix project crawler deep directory issue (hotfix).

[0.2.3] – 2015-12-09

Fixed

- Fix a few bugs related to project views.

[0.2.2] – 2015-11-30

Fixed

- Fix `SignacProjectCrawler.super()` bug.

[0.2.1] – 2015-11-29

Added

- Add support for Python 2.7
- Add signac-gui (early alpha)
- Allow specification of relative and default workspace paths
- Add the ability to create project views
- Add `Project.find_*` functions to search the workspace
- Add function to write and read state point hash tables

[0.2.0] – 2015-11-05

- Major consolidation of the package.
- Remove all hard dependencies, but six.

1.5 Support and Development

To get help using the **signac** package, join the [signac Slack workspace](#) or send an email to signac-support@umich.edu.

The **signac** package is hosted on [GitHub](#) and licensed under the open-source BSD 3-Clause license. Please use the [repository's issue tracker](#) to report bugs or request new features.

1.5.1 Code contributions

This project is open-source. Users are highly encouraged to contribute directly by implementing new features and fixing issues. Development for packages as part of the **signac** framework should follow the general development guidelines outlined [here](#).

A brief summary of contributing guidelines are outlined in the [CONTRIBUTING.md](#) file as part of the repository. All contributors must agree to the [Contributor Agreement](#) before their pull request can be merged.

Set up a development environment

Start by [forking](#) the project.

We highly recommend to setup a dedicated development environment, for example with [venv](#):

```
~ $ python -m venv ~/envs/signac-dev
~ $ source ~/envs/signac-dev/bin/activate
(signac-dev) ~ $ pip install pre-commit
```

or alternatively with [conda](#):

```
~ $ conda create -n signac-dev -c conda-forge python=3 pre-commit
~ $ conda activate signac-dev
```

Then clone your fork and install the package from source with:

```
(signac-dev) ~ $ cd path/to/my/fork/of/signac
(signac-dev) signac $ pip install -e .
```

The `-e` option stands for *editable*, which means that the package is directly loaded from the source code repository. That means any changes made to the source code are immediately reflected upon reloading the Python interpreter.

The [pre-commit](#) tool is used to enforce code style guidelines. To install the tool and configure pre-commit hooks, execute:

```
(signac-dev) signac $ pip install pre-commit
(signac-dev) signac $ pre-commit install
```

With the pre-commit hook, your code will be checked for syntax and style before you make a commit. The continuous integration pipeline for the package will perform these checks as well, so running these tests before committing / pushing will prevent the pipeline from failing due to style-related issues.

The development workflow

Prior to working on a patch, it is advisable to create an [issue](#) that describes the problem or proposed feature. This means that the code maintainers and other users get a chance to provide some input on the scope and possible limitations of the proposed changes, as well as advise on the actual implementation.

All code changes should be developed within a dedicated git branch and must all be related to each other. Unrelated changes, such as minor fixes to unrelated bugs encountered during implementation, spelling errors, and similar typographical mistakes must be developed within a separate branch.

Branches should be named after the following pattern: <prefix>/issue-<#>-optional-short-description. Choose from one of the following prefixes depending on the type of change:

- `fix/`: Any changes that fix the code and documentation.
- `feature/`: Any changes that introduce a new feature.
- `release/`: Reserved for release branches.

If your change does not seem to fall into any of the above mentioned categories, use `misc/`.

Once you are content with your changes, push the new branch to your forked repository and create a pull request into the main repository. Feel free to push a branch before completion to get input from the maintainers and other users, but make sure to add a comment that clarifies that the branch is not ready for merge yet.

Testing

Prior to fixing an issue, implement unit tests that *fail* for the described problem. New features must be tested with unit and integration tests. To run tests, execute:

```
(signac-dev) signac $ python -m pytest tests/
```

Building documentation

Building documentation requires the `sphinx` package which you will need to install into your development environment.

```
(signac-dev) signac $ pip install Sphinx sphinx_rtd_theme
```

Then you can build the documentation from within the `doc/` directory as part of the source code repository:

```
(signac-dev) signac $ cd doc/
(signac-dev) doc $ make html
```

Note: Documentation as part of the package should be largely limited to the API. More elaborate documentation on how to integrate **signac** into a computational workflow should be documented as part of the [framework documentation](#), which is maintained [here](#).

Updating the changelog

To update the changelog, add a one-line description to the `changelog.txt` file within the `next` section. For example:

```
next
----

- Fix issue with launching rockets to the moon.

[0.6.3] -- 2018-08-22
-----
```

(continues on next page)

(continued from previous page)

- Fix issue related to dynamic data spaces, ...

Just add the `next` section in case it doesn't exist yet.

CHAPTER 2

Indices and tables

- `genindex`
- `modindex`
- `search`

S

signac, 47
signac.cite, 53
signac.errors, 57
signac.sync, 54
signac.synced_collections.backends.collection_json,
65
signac.synced_collections.backends.collection_mongodb,
71
signac.synced_collections.backends.collection_redis,
74
signac.synced_collections.backends.collection_zarr,
76
signac.synced_collections.buffered.buffered_collection,
79
signac.synced_collections.buffered.file_buffered_collection,
80
signac.synced_collections.buffered.memory_buffered_collection,
83
signac.synced_collections.buffered.serialized_file_buffered_collection,
82
signac.synced_collections.data_types.synced_collection,
59
signac.synced_collections.data_types.synced_dict,
63
signac.synced_collections.data_types.synced_list,
64
signac.synced_collections.utils, 84
signac.synced_collections.validators,
86
signac.warnings, 57

Symbols

_load_and_save (class in `signac.synced_collections.data_types.synced_collection`), 63
 _load_save_type (attribute in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 61
 _backend (attribute in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 61
 _flush () (method in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 79
 _flush_buffer () (class method in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 80
 _from_base () (class method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 61
 _is_buffered (attribute in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 80
 _load () (method in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 80
 _load () (method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 61
 _load_from_buffer () (method in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 80
 _load_from_resource () (method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 62
 _register_validators () (class method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 62
 _save () (method in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 80
 _save () (method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 62
 _save_to_buffer () (method in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 80
 _save_to_resource () (method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 62
 _to_base () (method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 62
 update () (method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 62
 validate () (method in `signac.synced_collections.data_types.synced_collection.SyncedCollection`), 62
 AbstractTypeResolver (class in `signac.synced_collections.utils`), 84
 AbstractTypeResolver (class in `signac.synced_collections.utils`), 84
 always () (static method in `signac.sync.FileSync`), 55
 append () (method in `signac.synced_collections.data_types.synced_list.SyncedList`), 65
 AuthBackendError, 57
B
 backend_is_buffered () (method in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 80
 bibtex () (in module `signac.cite`), 53
 buffer_backend () (method in `signac.synced_collections.buffers.buffered_collection.BufferedCollection`), 80
 buffer_backend () (method in `signac.synced_collections.buffers.file_buffered_collection.FileBufferedCollection`), 81
 BufferedCollection (class in `signac.synced_collections.buffers.buffered_collection`), 79
 BufferedFileError, 58
 BufferedJSONDict (class in `signac.synced_collections.backends.collection_json`), 65
 BufferedJSONList (class in `signac.synced_collections.backends.collection_json`), 66
 BufferedJSONCollection (class in `signac.synced_collections.backends.collection_json`), 66

BufferedJSONDict (class in disable_multithreading()
signac.synced_collections.backends.collection_json),
 66 (signac.synced_collections.buffers.file_buffered_collection.FileBufferedJSONDict class method), 81

BufferedJSONList (class in disable_multithreading()
signac.synced_collections.backends.collection_json),
 66 (signac.synced_collections.data_types.synced_collection.SyncedCollection class method), 63

BufferException, 58 doc (signac.contrib.job.Job attribute), 34

build_job_search_index() (signac.Project doc (signac.Project attribute), 18
 method), 15 DocSync (class in signac.sync), 55

build_job_statepoint_index() DocSync.ByKey (class in signac.sync), 55
 (signac.Project method), 16 document (signac.contrib.job.Job attribute), 34
 document (signac.Project attribute), 18
 DocumentSyncConflict, 58
 dump() (signac.Collection method), 39
 dump_statepoints() (signac.Project method), 19

C

check() (signac.Project method), 16

clear() (signac.Collection method), 39

clear() (signac.contrib.job.Job method), 33

clear() (signac.H5Store method), 45

clear() (signac.synced_collections.data_types.synced_dict.SyncedDict
 method), 64 (signac.synced_collections.buffers.file_buffered_collection.FileBufferedDict class method), 81

clear() (signac.synced_collections.data_types.synced_list.SyncedList
 method), 65 enable_multithreading()

client (signac.synced_collections.backends.collection_redis.RedisCollection
 attribute), 75 (signac.synced_collections.data_types.synced_collection.SyncedCollection class method), 63

clone() (signac.Project method), 16

close() (signac.Collection method), 39

close() (signac.contrib.job.Job method), 34

close() (signac.H5Store method), 45

cls (signac.H5StoreManager attribute), 47

codec (signac.synced_collections.backends.collection_zarr.ZarrCollection
 attribute), 77 Error, 58

Collection (class in signac), 38 export() (in module signac), 50

collection (signac.synced_collections.backends.collection_mongodb.MongoDBCollection
 attribute), 72 export_one() (in module signac), 50

config (signac.Project attribute), 17 export_pymongo() (in module signac), 51

ConfigError, 58 export_to() (signac.contrib.project.JobsCursor
 method), 30

COPY (signac.sync.DocSync attribute), 55 export_to() (signac.Project method), 19

create_access_module() (signac.Project export_to_mirror() (in module signac), 51
 method), 17 extend() (signac.synced_collections.data_types.synced_list.SyncedList
 method), 65

create_linked_view() (signac.Project method), 17

D

data (signac.contrib.job.Job attribute), 34

data (signac.Project attribute), 18

default() (in module signac.synced_collections.utils), 86

default() (signac.synced_collections.utils.SyncedCollectionJSONEncoder
 method), 86

delete_many() (signac.Collection method), 39

delete_one() (signac.Collection method), 39

DestinationExistsError, 58

detect_schema() (signac.Project method), 18

diff_jobs() (in module signac), 48

E

enable_multithreading()

Error, 58

export() (in module signac), 50

export_one() (in module signac), 50

export_pymongo() (in module signac), 51

export_to() (signac.contrib.project.JobsCursor
 method), 30

export_to() (signac.Project method), 19

export_to_mirror() (in module signac), 51

extend() (signac.synced_collections.data_types.synced_list.SyncedList
 method), 65

F

fetch() (in module signac), 49

FetchError, 58

file (signac.H5Store attribute), 45

FileBufferedCollection (class in
 signac.synced_collections.buffers.file_buffered_collection),
 80

filename (signac.errors.FileSyncConflict attribute), 58

filename (signac.H5Store attribute), 46

filename (signac.synced_collections.backends.collection_json.JSONCollection
 attribute), 68

FileSync (class in signac.sync), 55

FileSync.Ask (class in signac.sync), 55

FileSyncConflict, 58

find() (signac.Collection method), 39

find_job_ids() (signac.Project method), 20

find_jobs() (signac.Project method), 20

find_one() (signac.Collection method), 41

flush() (in module signac), 53
 flush() (signac.Collection method), 41
 flush() (signac.H5Store method), 46
 fn() (signac.contrib.job.Job method), 34
 fn() (signac.Project method), 21
 FN_CACHE (signac.Project attribute), 15
 FN_DOCUMENT (signac.contrib.job.Job attribute), 33
 FN_DOCUMENT (signac.Project attribute), 15
 FN_MANIFEST (signac.contrib.job.Job attribute), 33
 FN_STATEPOINTS (signac.Project attribute), 15

G

get() (signac.H5Store method), 46
 get() (signac.synced_collections.data_types.synced_dict.SyncedDict class method), 64
 get_buffer_capacity() (signac.synced_collections.buffers.file_buffered_collection.FileBufferedCollection class method), 81
 get_current_buffer_size() (signac.synced_collections.buffers.file_buffered_collection.FileBufferedCollection class method), 82
 get_database() (in module signac), 49
 get_id() (signac.contrib.job.Job method), 34
 get_id() (signac.Project method), 21
 get_job() (in module signac), 48
 get_job() (signac.Project class method), 21
 get_project() (in module signac), 47
 get_project() (signac.Project class method), 21
 get_statepoint() (signac.Project method), 21
 get_type() (signac.synced_collections.utils.AbstractTypeResolver class method), 85
 group (signac.synced_collections.backends.collection_zarr.ZarrCollection attribute), 77
 groupby() (signac.contrib.project.JobsCursor method), 30
 groupby() (signac.Project method), 22
 groupbydoc() (signac.contrib.project.JobsCursor method), 31
 groupbydoc() (signac.Project method), 23

H

H5Store (class in signac), 44
 H5StoreManager (class in signac), 46
 HDF5 integration, 4

I

id (signac.contrib.job.Job attribute), 35
 id (signac.Project attribute), 23
 ids (signac.Collection attribute), 41
 import_from() (signac.Project method), 23
 IncompatibleSchemaVersion, 58
 index() (in module signac), 52
 index() (signac.Collection method), 42
 index() (signac.Project method), 24

index_files() (in module signac), 52
 init() (signac.contrib.job.Job method), 35
 init_project() (in module signac), 48
 init_project() (signac.Project class method), 25
 insert() (signac.synced_collections.data_types.synced_list.SyncedList method), 65
 insert_one() (signac.Collection method), 42
 InvalidKeyError, 58
 is_base_type() (signac.synced_collections.data_types.synced_collection.SyncedCollection class method), 63
 is_base_type() (signac.synced_collections.data_types.synced_dict.SyncedDict class method), 64
 is_base_type() (signac.synced_collections.data_types.synced_list.SyncedList class method), 65
 isfile() (signac.contrib.job.Job method), 35
 isfile() (signac.Project method), 25
 items() (signac.synced_collections.data_types.synced_dict.SyncedDict method), 64
 Job (class in signac.contrib.job), 33
 JobsCorruptedError, 59
 JobsCursor (class in signac.contrib.project), 30
 json_attr_dict_validator() (in module signac.synced_collections.backends.collection_json), 71
 json_format_validator() (in module signac.synced_collections.validators), 86
 JSONAttrDict (class in signac.synced_collections.backends.collection_json), 66
 JSONAttrList (class in signac.synced_collections.backends.collection_json), 67
 JSONCollection (class in signac.synced_collections.backends.collection_json), 68
 JSONDict (class in signac.synced_collections.backends.collection_json), 68
 JSONDict (in module signac), 44
 JSONList (class in signac.synced_collections.backends.collection_json), 69

K

key (signac.synced_collections.backends.collection_redis.RedisCollection attribute), 75
 KEY_DATA (signac.contrib.job.Job attribute), 33
 KEY_DATA (signac.Project attribute), 15
 keys (signac.errors.DocumentSyncConflict attribute), 58
 keys() (signac.H5Store method), 46
 keys() (signac.H5StoreManager method), 47
 keys() (signac.sync.FileSync class method), 55

keys () (*signac.synced_collections.data_types.synced_dict.SyncedDict* method), 64

KeyTypeError, 59

M

main () (*signac.Collection* method), 42

MemoryBufferedJSONAttrDict (class in *signac.synced_collections.backends.collection_json*), 70

MemoryBufferedJSONAttrList (class in *signac.synced_collections.backends.collection_json*), 70

MemoryBufferedJSONCollection (class in *signac.synced_collections.backends.collection_json*), 70

MemoryBufferedJSONDict (class in *signac.synced_collections.backends.collection_json*), 71

MemoryBufferedJSONList (class in *signac.synced_collections.backends.collection_json*), 71

min_len_unique_id () (*signac.Project* method), 25

mode (*signac.H5Store* attribute), 46

MongoDB database backend, 4

MongoDBCollection (class in *signac.synced_collections.backends.collection_mongodb*), 71

MongoBDDict (class in *signac.synced_collections.backends.collection_mongodb*), 72

MongoDBList (class in *signac.synced_collections.backends.collection_mongodb*), 73

move () (*signac.contrib.job.Job* method), 35

N

name (*signac.synced_collections.backends.collection_zarr.ZarrCollection* attribute), 77

never () (*signac.sync.FileSync* static method), 55

next () (*signac.contrib.project.JobsCursor* method), 32

no_dot_in_key () (in module *signac.synced_collections.validators*), 87

NO_SYNC (*signac.sync.DocSync* attribute), 55

num_jobs () (*signac.Project* method), 25

O

open () (*signac.Collection* class method), 43

open () (*signac.contrib.job.Job* method), 35

open () (*signac.H5Store* method), 46

open_job () (*signac.Project* method), 25

P

pop () (*signac.H5Store* method), 46

popitem () (*signac.H5Store* method), 46

popitem () (*signac.synced_collections.data_types.synced_dict.SyncedDict* method), 64

prefix (*signac.H5StoreManager* attribute), 47

primary_key (*signac.Collection* attribute), 43

Project (class in *signac*), 15

R

read_json () (*signac.Collection* class method), 43

read_statepoints () (*signac.Project* method), 26

RedisCollection (class in *signac.synced_collections.backends.collection_redis*), 74

RedisDict (class in *signac.synced_collections.backends.collection_redis*), 75

RedisList (class in *signac.synced_collections.backends.collection_redis*), 75

reference () (in module *signac.cite*), 53

remove () (*signac.contrib.job.Job* method), 36

remove () (*signac.synced_collections.data_types.synced_list.SyncedList* method), 65

remove () (*signac.Project* method), 26

replace_one () (*signac.Collection* method), 44

require_string_key () (in module *signac.synced_collections.validators*), 87

reset () (*signac.contrib.job.Job* method), 36

reset () (*signac.synced_collections.data_types.synced_dict.SyncedDict* method), 64

reset () (*signac.synced_collections.data_types.synced_list.SyncedList* method), 65

reset_statepoint () (*signac.contrib.job.Job* method), 36

reset_statepoint () (*signac.Project* method), 26

root_directory () (*signac.Project* method), 27

S

SchemaSyncConflict, 59

SerializedFileBufferedCollection (class in *signac.synced_collections.buffers.serialized_file_buffered_collection*), 82

set_buffer_capacity () (*signac.synced_collections.buffers.file_buffered_collection.FileBufferedCollection* class method), 82

setdefault () (*signac.H5Store* method), 46

setdefault () (*signac.synced_collections.data_types.synced_dict.SyncedDict* method), 64

SharedMemoryFileBufferedCollection (class in *signac.synced_collections.buffers.memory_buffered_collection*), 83

signac (module), 47

- signac.cite (*module*), 53
- signac.errors (*module*), 57
- signac.sync (*module*), 54
- signac.synced_collections.backends.collection_zarr (*module*), 65
- signac.synced_collections.backends.collection_mongodb (*module*), 71
- signac.synced_collections.backends.collection_redis (*module*), 74
- signac.synced_collections.backends.collection_zarr (*module*), 76
- signac.synced_collections.buffered_collection (*module*), 79
- signac.synced_collections.buffered_file_buffered_collection (*module*), 80
- signac.synced_collections.buffered_memory_buffered_collection (*module*), 83
- signac.synced_collections.buffered_serialized_file_buffered_collection (*module*), 82
- signac.synced_collections.data_types.synced_collection (*module*), 59
- signac.synced_collections.data_types.synced_dict (*module*), 63
- signac.synced_collections.data_types.synced_list (*module*), 64
- signac.synced_collections.utils (*module*), 84
- signac.synced_collections.validators (*module*), 86
- signac.warnings (*module*), 57
- SignacDeprecationWarning, 57
- sp (*signac.contrib.job.Job* attribute), 36
- statepoint (*signac.contrib.job.Job* attribute), 36
- StatepointParsingError, 59
- stores (*signac.contrib.job.Job* attribute), 36
- stores (*signac.Project* attribute), 27
- sync () (*signac.contrib.job.Job* method), 37
- sync () (*signac.Project* method), 27
- sync_jobs () (*in module signac.sync*), 55
- sync_projects () (*in module signac.sync*), 56
- SyncConflict, 59
- SyncedCollection (class in *signac.synced_collections.data_types.synced_collection*), 59
- SyncedCollectionJSONEncoder (class in *signac.synced_collections.utils*), 85
- SyncedDict (class in *signac.synced_collections.data_types.synced_dict*), 63
- SyncedList (class in *signac.synced_collections.data_types.synced_list*), 64
- T**
- temporary_project () (*signac.Project* method), 28
- TemporaryProject () (*in module signac*), 47
- to_dataframe () (*signac.contrib.project.JobsCursor* method), 32
- to_dataframe () (*signac.Project* method), 28
- to_json () (*signac.Collection* method), 44
- type_map (*signac.synced_collections.utils.AbstractTypeResolver* attribute), 85
- U**
- uid (*signac.synced_collections.backends.collection_mongodb.MongoDBC* attribute), 72
- update () (*signac.Collection* method), 44
- update () (*signac.H5Store* method), 46
- update () (*signac.sync.DocSync* static method), 55
- update () (*signac.sync.FileSync* static method), 55
- update () (*signac.synced_collections.data_types.synced_dict.SyncedDict* method), 64
- update_cache () (*signac.Project* method), 28
- update_statepoint () (*signac.contrib.job.Job* method), 37
- update_statepoint () (*signac.Project* method), 28
- V**
- values () (*signac.H5Store* method), 46
- values () (*signac.synced_collections.data_types.synced_dict.SyncedDict* method), 64
- W**
- workspace () (*signac.contrib.job.Job* method), 38
- workspace () (*signac.Project* method), 29
- WorkspaceError, 59
- write_statepoints () (*signac.Project* method), 29
- ws (*signac.contrib.job.Job* attribute), 38
- Z**
- ZarrCollection (class in *signac.synced_collections.backends.collection_zarr*), 76
- ZarrDict (class in *signac.synced_collections.backends.collection_zarr*), 77
- ZarrList (class in *signac.synced_collections.backends.collection_zarr*), 78