

# Package ‘rslurm’

August 29, 2016

**Type** Package

**Title** Submit R Calculations to a 'SLURM' Cluster

**Description** Functions that simplify the R interface the 'SLURM' cluster workload manager, and automate the process of dividing a parallel calculation across cluster nodes.

**Version** 0.3.1

**License** GPL-3

**URL** <https://github.com/SESYNC-ci/rslurm>

**BugReports** <https://github.com/SESYNC-ci/rslurm/issues>

**Depends** R (>= 3.2.0)

**Imports** parallel, whisker (>= 0.3)

**RoxygenNote** 5.0.1

**Suggests** testthat, knitr, rmarkdown

**VignetteBuilder** knitr

**NeedsCompilation** no

**Author** Philippe Marchand [aut, cre],  
Mike Smorul [ctb]

**Maintainer** Philippe Marchand <pmarchand@sesync.org>

**Repository** CRAN

**Date/Publication** 2016-06-18 19:05:41

## R topics documented:

rslurm-package	2
cancel_slurm	3
cleanup_files	3
get_slurm_out	4
print_job_status	5
slurm_apply	6
slurm_call	7
slurm_job	9

---

 rslurm-package

*rslurm: Submit R calculations to a SLURM cluster*


---

## Description

This package automates the process of sending simple function calls or parallel calculations to a cluster using the SLURM workload manager.

## Overview

This package includes two core functions used to send computations to a SLURM cluster. While `slurm_call` executes a function using a single set of parameters (passed as a list), `slurm_apply` evaluates the function in parallel for multiple sets of parameters grouped in a data frame. `slurm_apply` automatically splits the parameter sets into equal-size chunks, each chunk to be processed by a separate cluster node. It uses functions from the `parallel` package to parallelize computations within each node.

The output of `slurm_apply` or `slurm_call` is a `slurm_job` object that serves as an input to the other functions in the package: `print_job_status`, `cancel_slurm`, `get_slurm_out` and `cleanup_files`.

For bug reports or questions about this package, contact Philippe Marchand ([pmarchand@sesync.org](mailto:pmarchand@sesync.org)).

## Function Specification

To be compatible with `slurm_apply`, a function may accept any number of single value parameters. The names of these parameters must match the column names of the `params` data frame supplied. There are no restrictions on the types of parameters passed as a list to `slurm_call`.

If the function passed to `slurm_call` or `slurm_apply` requires knowledge of any R objects (data, custom helper functions) besides `params`, a character vector corresponding to their names should be passed to the optional `add_objects` argument.

When parallelizing a function, since any error will interrupt all calculations for the current node, it may be useful to wrap expressions which may generate errors into a `try` or `tryCatch` function. This will ensure the computation continues with the next parameter set after reporting the error.

## Output Format

The default output format for `get_slurm_out` (`outtype = "raw"`) is a list where each element is the return value of one function call. If the function passed to `slurm_apply` produces a vector output, you may use `outtype = "table"` to collect the output in a single data frame, with one row by function call.

## Examples

```
## Not run:
# Create a data frame of mean/sd values for normal distributions
pars <- data.frame(par_m = seq(-10, 10, length.out = 1000),
                  par_sd = seq(0.1, 10, length.out = 1000))
```

```
# Create a function to parallelize
fctest <- function(par_m, par_sd) {
  samp <- rnorm(10^7, par_m, par_sd)
  c(s_m = mean(samp), s_sd = sd(samp))
}

sjob1 <- slurm_apply(fctest, pars)
print_job_status(sjob1)
res <- get_slurm_out(sjob1, "table")
all.equal(pars, res) # Confirm correct output
cleanup_files(sjob1)

## End(Not run)
```

---

cancel_slurm	<i>Cancels a scheduled SLURM job</i>
--------------	--------------------------------------

---

### Description

This function cancels the specified SLURM job by invoking the SLURM `scancel` command. It does *not* delete the temporary files (e.g. scripts) created by [slurm\\_apply](#) or [slurm\\_call](#). Use [cleanup\\_files](#) to remove those files.

### Usage

```
cancel_slurm(slr_job)
```

### Arguments

`slr_job` A `slurm_job` object.

### See Also

[cleanup\\_files](#)

---

cleanup_files	<i>Deletes temporary files associated with a SLURM job</i>
---------------	--

---

### Description

This function deletes all temporary files associated with the specified SLURM job, including files created by [slurm\\_apply](#) or [slurm\\_call](#), as well as outputs from the cluster. These files should be located in the `._rslurm_[jobname]` folder of the current working directory.

**Usage**

```
cleanup_files(slr_job)
```

**Arguments**

slr\_job            A slurm\_job object.

**See Also**

[slurm\\_apply](#), [slurm\\_call](#)

**Examples**

```
## Not run:
sjob <- slurm_apply(func, pars)
print_job_status(sjob) # Prints console/error output once job is completed.
func_result <- get_slurm_out(sjob, "table") # Loads output data into R.
cleanup_files(sjob)

## End(Not run)
```

---

get\_slurm\_out

*Reads the output of a function calculated on the SLURM cluster*

---

**Description**

This function reads all function output files (one by cluster node used) from the specified SLURM job and returns the result in a single data frame (if "table" format selected) or list (if "raw" format selected). It doesn't record any messages (including warnings or errors) output to the R console during the computation; these can be consulted by invoking [print\\_job\\_status](#).

**Usage**

```
get_slurm_out(slr_job, outtype = "raw")
```

**Arguments**

slr\_job            A slurm\_job object.

outtype           Can be "table" or "raw", see "Value" below for details.

**Details**

The outtype option is only relevant for jobs submitted with `slurm_apply`. Jobs sent with `slurm_call` only return a single object, and setting `outtype = "table"` creates an error in that case.

**Value**

If `outtype = "table"`: A data frame with one column by return value of the function passed to `slurm_apply`, where each row is the output of the corresponding row in the `params` data frame passed to `slurm_apply`.

If `outtype = "raw"`: A list where each element is the output of the function passed to `slurm_apply` for the corresponding row in the `params` data frame passed to `slurm_apply`.

**See Also**

[slurm\\_apply](#), [slurm\\_call](#)

---

<code>print_job_status</code>	<i>Prints the status of a SLURM job and, if completed, its console/error output</i>
-------------------------------	---

---

**Description**

Prints the status of a SLURM job and, if completed, its console/error output.

**Usage**

```
print_job_status(slr_job)
```

**Arguments**

`slr_job`      A `slurm_job` object.

**Details**

If the specified SLURM job is still in the queue or running, this function prints its current status (as output by the SLURM `squeue` command). The output displays one row by node currently running part of the job ("R" in the "ST" column) and how long it has been running ("TIME"). One row indicates the portions of the job still in queue ("PD" in the "ST" column), if any.

If all portions of the job have completed or stopped, the function prints the console and error output, if any, generated by each node.

---

 slurm\_apply

*Parallel execution of a function on the SLURM cluster*


---

### Description

Use `slurm_apply` to compute function over multiple sets of parameters in parallel, spread across multiple nodes of a SLURM cluster.

### Usage

```
slurm_apply(f, params, jobname = NA, nodes = 2, cpus_per_node = 2,
  add_objects = NULL, pkgs = rev(.packages()), slurm_options = list(),
  submit = TRUE)
```

### Arguments

<code>f</code>	A function that accepts one or many single values as parameters and may return any type of R object.
<code>params</code>	A data frame of parameter values to apply <code>f</code> to. Each column corresponds to a parameter of <code>f</code> ( <i>Note</i> : names must match) and each row corresponds to a separate function call.
<code>jobname</code>	The name of the SLURM job; if NA, it is assigned a random name of the form "slr####".
<code>nodes</code>	The (maximum) number of cluster nodes to spread the calculation over. <code>slurm_apply</code> automatically divides <code>params</code> in chunks of approximately equal size to send to each node. Less nodes are allocated if the parameter set is too small to use all CPUs on the requested nodes.
<code>cpus_per_node</code>	The number of CPUs per node on the cluster; determines how many processes are run in parallel per node.
<code>add_objects</code>	A character vector containing the name of R objects to be saved in a .RData file and loaded on each cluster node prior to calling <code>f</code> .
<code>pkgs</code>	A character vector containing the names of packages that must be loaded on each cluster node. By default, it includes all packages loaded by the user when <code>slurm_apply</code> is called.
<code>slurm_options</code>	A named list of options recognized by <code>sbatch</code> ; see Details below for more information.
<code>submit</code>	Whether or not to submit the job to the cluster with <code>sbatch</code> ; see Details below for more information.

### Details

This function creates a temporary folder ("`_rslurm_[jobname]`") in the current directory, holding the .RData files, the R script and the Bash submission script generated for the SLURM job.

The set of input parameters is divided in equal chunks sent to each node, and `f` is evaluated in parallel within each node using functions from the `parallel` R package. The names of any other R objects (besides `params`) that `f` needs to access should be included in `add_objects`.

Use `slurm_options` to set any option recognized by `sbatch`, e.g. `slurm_options = list(time = "1:00:00", share = TRUE)`. See <http://slurm.schedmd.com/sbatch.html> for details on possible options. Note that full names must be used (e.g. "time" rather than "t") and that flags (such as "share") must be specified as TRUE. The "array", "job-name", "nodes" and "output" options are already determined by `slurm_apply` and should not be manually set.

When processing the computation job, the SLURM cluster will output two types of files in the temporary folder: those containing the return values of the function for each subset of parameters ("results\_[node\_id].RData") and those containing any console or error output produced by R on each node ("slurm\_[node\_id].out").

If `submit = TRUE`, the job is sent to the cluster and a confirmation message (or error) is output to the console. If `submit = FALSE`, a message indicates the location of the saved data and script files; the job can be submitted manually by running the shell command `sbatch submit.sh` from that directory.

After sending the job to the SLURM cluster, `slurm_apply` returns a `slurm_job` object which can be used to cancel the job, get the job status or output, and delete the temporary files associated with it. See the description of the related functions for more details.

### Value

A `slurm_job` object containing the jobname and the number of nodes effectively used.

### See Also

[slurm\\_call](#) to evaluate a single function call.

[cancel\\_slurm](#), [cleanup\\_files](#), [get\\_slurm\\_out](#) and [print\\_job\\_status](#) which use the output of this function.

### Examples

```
## Not run:
sjob <- slurm_apply(func, pars)
print_job_status(sjob) # Prints console/error output once job is completed.
func_result <- get_slurm_out(sjob, "table") # Loads output data into R.
cleanup_files(sjob)

## End(Not run)
```

---

slurm\_call

*Execution of a single function call on the SLURM cluster*

---

### Description

Use `slurm_call` to perform a single function evaluation on the SLURM cluster.

**Usage**

```
slurm_call(f, params, jobname = NA, add_objects = NULL,
           pkgs = rev(.packages()), slurm_options = list(), submit = TRUE)
```

**Arguments**

<code>f</code>	Any R function.
<code>params</code>	A named list of parameters to pass to <code>f</code> .
<code>jobname</code>	The name of the SLURM job; if <code>NA</code> , it is assigned a random name of the form "slr####".
<code>add_objects</code>	A character vector containing the name of R objects to be saved in a .RData file and loaded on each cluster node prior to calling <code>f</code> .
<code>pkgs</code>	A character vector containing the names of packages that must be loaded on each cluster node. By default, it includes all packages loaded by the user when <code>slurm_call</code> is called.
<code>slurm_options</code>	A named list of options recognized by <code>sbatch</code> ; see Details below for more information.
<code>submit</code>	Whether or not to submit the job to the cluster with <code>sbatch</code> ; see Details below for more information.

**Details**

This function creates a temporary folder ("`_rslurm_[jobname]`") in the current directory, holding the .RData files, the R script and the Bash submission script generated for the SLURM job.

The names of any other R objects (besides `params`) that `f` needs to access should be listed in the `add_objects` argument.

Use `slurm_options` to set any option recognized by `sbatch`, e.g. `slurm_options = list(time = "1:00:00", share = TRUE)`. See <http://slurm.schedmd.com/sbatch.html> for details on possible options. Note that full names must be used (e.g. "time" rather than "t") and that flags (such as "share") must be specified as `TRUE`. The "job-name", "ntasks" and "output" options are already determined by `slurm_call` and should not be manually set.

When processing the computation job, the SLURM cluster will output two files in the temporary folder: one with the return value of the function ("`results_0.RData`") and one containing any console or error output produced by R ("`slurm_[node_id].out`").

If `submit = TRUE`, the job is sent to the cluster and a confirmation message (or error) is output to the console. If `submit = FALSE`, a message indicates the location of the saved data and script files; the job can be submitted manually by running the shell command `sbatch submit.sh` from that directory.

After sending the job to the SLURM cluster, `slurm_call` returns a `slurm_job` object which can be used to cancel the job, get the job status or output, and delete the temporary files associated with it. See the description of the related functions for more details.

**Value**

A `slurm_job` object containing the `jobname` and the number of nodes effectively used.



**See Also**

[slurm\\_apply](#) to parallelize a function over a parameter set.

[cancel\\_slurm](#), [cleanup\\_files](#), [get\\_slurm\\_out](#) and [print\\_job\\_status](#) which use the output of this function.

---

slurm_job	<i>Create a slurm_job object</i>
-----------	----------------------------------

---

**Description**

This function creates a `slurm_job` object which can be passed to other functions such as [cancel\\_slurm](#), [cleanup\\_files](#), [get\\_slurm\\_out](#) and [print\\_job\\_status](#).

**Usage**

```
slurm_job(jobname, nodes)
```

**Arguments**

jobname	The name of the SLURM job. The rslurm-generated scripts and output files associated with a job should be found in the <code>_rslurm_[jobname]</code> folder.
nodes	The number of cluster nodes used by that job.

**Details**

In general, `slurm_job` objects are created automatically as the output of [slurm\\_apply](#) or [slurm\\_call](#), but it may be necessary to manually recreate one if the job was submitted in a different R session.

**Value**

A `slurm_job` object.

# Index

cancel\_slurm, [2](#), [3](#), [7](#), [9](#)  
cleanup\_files, [2](#), [3](#), [3](#), [7](#), [9](#)  
  
get\_slurm\_out, [2](#), [4](#), [7](#), [9](#)  
  
parallel, [2](#)  
print\_job\_status, [2](#), [4](#), [5](#), [7](#), [9](#)  
  
rslurm (rslurm-package), [2](#)  
rslurm-package, [2](#)  
  
slurm\_apply, [2-5](#), [6](#), [9](#)  
slurm\_call, [2-5](#), [7](#), [7](#), [9](#)  
slurm\_job, [9](#)  
  
try, [2](#)  
tryCatch, [2](#)