

Package ‘doFuture’

October 20, 2017

Version 0.6.0

Title A Universal Foreach Parallel Adaptor using the Future API of the 'future' Package

Depends future (>= 1.6.2), foreach (>= 1.4.3), iterators (>= 1.0.8), parallel, utils

Suggests doRNG (>= 1.6.6), markdown, R.rsp

VignetteBuilder R.rsp

Description Provides a '%dopar%' adaptor such that any type of futures can be used as backends for the 'foreach' framework.

License LGPL (>= 2.1)

LazyLoad TRUE

URL <https://github.com/HenrikBengtsson/doFuture>

BugReports <https://github.com/HenrikBengtsson/doFuture/issues>

RoxygenNote 6.0.1

NeedsCompilation no

Author Henrik Bengtsson [aut, cre, cph]

Maintainer Henrik Bengtsson <henrikb@braju.com>

Repository CRAN

Date/Publication 2017-10-20 10:07:31 UTC

R topics documented:

doFuture	2
registerDoFuture	4
Index	5

Description

The **doFuture** package provides a `%dopar%` adaptor for the **foreach** package such that *any* type of future (that is supported by Future API of the **future** package) can be used for asynchronous (parallel/distributed) or synchronous (sequential) processing.

Details

In other words, *if a computational backend is supported via the Future API, it'll be automatically available for all functions and packages making using the **foreach** framework*. Neither the developer nor the end user has to change any code.

Usage

To use futures with the **foreach** package, load **doFuture**, use `registerDoFuture()` to register it to be used as a `%dopar%` adaptor (need to ever use `%do%`). After this, how and where the computations are performed is controlled solely by the future strategy set, which is controlled by `future::plan()`. For example:

- `plan(multiprocess)`: multiple R processes on the local machine.
- `plan(cluster, workers = c("n1", "n2", "n2", "n3"))`: multiple R processes on external machines.

See the **future** package for more examples.

Built-in backends

The built-in backends of **doFuture** are for instance multicore (forked processes), multisession (background R sessions), and ad-hoc cluster (background R sessions on local and / or remote machines). Additional futures are provided by other "future" packages (see below for some examples).

Backends for high-performance compute clusters

The **future.BatchJobs** package provides support for high-performance compute (HPC) cluster schedulers such as SGE, Slurm, and TORQUE / PBS. For example,

- `plan(batchjobs_slurm)`: Process via a Slurm scheduler job queue.
- `plan(batchjobs_torque)`: Process via a TORQUE / PBS scheduler job queue.

This builds on top of the queuing framework that the **BatchJobs** package provides. For more details on backend configuration, please see the **future.BatchJobs** and **BatchJobs** packages.

Global variables and packages

Unless running locally in the global environment (= at the R prompt), the **foreach** package requires you do specify what global variables and packages need to be available and attached in order for the "foreach" expression to be evaluated properly. It is not uncommon to get errors on one or missing variables when moving from running a `res <- foreach() %dopar% { ... }` statement on the local machine to, say, another machine on the same network. The solution to the problem is to explicitly export those variables by specifying them in the `.export` argument to `foreach()`, e.g. `foreach(..., .export = c("mu", "sigma"))`. Likewise, if the expression needs specific packages to be attached, they can be listed in argument `.packages` of `foreach()`.

When using `doFuture::registerDoFuture()`, the above becomes less critical, because by default the Future API identifies all globals and all packages automatically (via static code inspection). This is done exactly the same way regardless of future backend. This automatic identification of globals and packages is illustrated by the below example, which does *not* specify `.export = c("my_stat")`. This works because the future framework detects that function `my_stat()` is needed and makes sure it is exported. If you would use, say, `cl <- parallel::makeCluster(2)` and `doParallel::registerDoParallel(cl)`, you would get a run-time error on 'Error in { : task 1 failed - "could not find function "my_stat"'.

Having said this, note that, in order for your "foreach" code to work everywhere and with other types of foreach adaptors as well, you may want to make sure that you always specify arguments `.export` and `.packages`.

Load balancing ("chunking")

Whether load balancing ("chunking") should take place or not can be controlled by `foreach()` option `.options.future = list(scheduling = <value>)`. The value `scheduling` specifies the average number of futures ("chunks") per worker. If `0.0`, then a single future is used to process all iterations - none of the other workers are not used. If `1.0` or `TRUE`, then one future per worker is used. If `2.0`, then each worker will process two futures (if there are enough iterations). If `Inf` or `FALSE`, then one future per iteration is used. The default value is `scheduling = 1.0`, unless option `.options.multicore = list(preschedule = <logical>)` is set, which in case that becomes the default. In other words, it is also possible to disable load balancing by using `.options.multicore = list(preschedule = FALSE)`.

Examples

```
library("doFuture")
registerDoFuture()
plan(multiprocess)

my_stat <- function(x) {
  median(x)
}

my_experiment <- function(n, mu = 0.0, sigma = 1.0) {
  foreach(i = 1:n) %dopar% {
    x <- rnorm(i, mean = mu, sd = sigma)
    list(mu = mean(x), sigma = sd(x), own = my_stat(x))
  }
}
```

```
}  
  
y <- my_experiment(n = 3)  
str(y)
```

registerDoFuture	<i>Registers the future %dopar% backend</i>
------------------	---

Description

Register the [doFuture](#) parallel adaptor to be used by the **foreach** package.

Usage

```
registerDoFuture()
```

Value

Nothing

Index

*Topic **utilities**

registerDoFuture, 4

doFuture, 2, 4

doFuture-package (doFuture), 2

foreach, 3

future::plan, 2

registerDoFuture, 2, 4