

Package ‘future’

July 1, 2023

Version 1.33.0

Title Unified Parallel and Distributed Processing in R for Everyone

Imports digest, globals (>= 0.16.1), listenv (>= 0.8.0), parallel,
parallelly (>= 1.34.0), utils

Suggests methods, RnpcBLASctl, R.rsp, markdown

VignetteBuilder R.rsp

Description The purpose of this package is to provide a lightweight and unified Future API for sequential and parallel processing of R expression via futures. The simplest way to evaluate an expression in parallel is to use `x %<-% { expression }` with `plan(multisession)`. This package implements sequential, multicore, multisession, and cluster futures. With these, R expressions can be evaluated on the local machine, in parallel a set of local machines, or distributed on a mix of local and remote machines.

Extensions to this package implement additional backends for processing futures via compute cluster schedulers, etc.

Because of its unified API, there is no need to modify any code in order to switch from sequential on the local machine to, say, distributed processing on a remote compute cluster.

Another strength of this package is that global variables and functions are automatically identified and exported as needed, making it straightforward to tweak existing code to make use of futures.

License LGPL (>= 2.1)

LazyLoad TRUE

ByteCompile TRUE

URL <https://future.futureverse.org>,
<https://github.com/HenrikBengtsson/future>

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

RoxygenNote 7.2.3

NeedsCompilation no

Author Henrik Bengtsson [aut, cre, cph]

Maintainer Henrik Bengtsson <henrikb@braju.com>

Repository CRAN

Date/Publication 2023-07-01 18:00:03 UTC

R topics documented:

backtrace	2
cluster	3
clusterExportSticky	5
future	6
future.options	12
futureOf	16
futures	17
futureSessionInfo	18
multicore	18
multisession	20
nbrOfWorkers	22
plan	23
re-exports	26
resolve	27
resolved	28
sequential	29
tweak	30
value	30
%conditions%	31
%globals%	32
%label%	32
%lazy%	33
%plan%	33
%seed%	34
%stdout%	34
%tweak%	35
Index	36

backtrace	<i>Back trace the expressions evaluated when an error was caught</i>
-----------	--

Description

Back trace the expressions evaluated when an error was caught

Usage

```
backtrace(future, envir = parent.frame(), ...)
```

Arguments

future	A future with a caught error.
envir	the environment where to locate the future.
...	Not used.

Value

A list with the future's call stack that led up to the error.

Examples

```
my_log <- function(x) log(x)
foo <- function(...) my_log(...)

f <- future({ foo("a") })
res <- tryCatch({
  v <- value(f)
}, error = function(ex) {
  t <- backtrace(f)
  print(t)
})
```

cluster	<i>Create a cluster future whose value will be resolved asynchronously in a parallel process</i>
---------	--

Description

A cluster future is a future that uses cluster evaluation, which means that its *value is computed and resolved in parallel in another process*.

Usage

```
cluster(
  ...,
  persistent = FALSE,
  workers = availableWorkers(),
  envir = parent.frame()
)
```

Arguments

...	Additional named elements passed to <code>ClusterFuture()</code> .
persistent	If FALSE, the evaluation environment is cleared from objects prior to the evaluation of the future.

workers	A cluster object, a character vector of host names, a positive numeric scalar, or a function. If a character vector or a numeric scalar, a <code>cluster</code> object is created using <code>makeClusterPSOCK(workers)</code> . If a function, it is called without arguments <i>when the future is created</i> and its value is used to configure the workers. The function should return any of the above types.
envir	The environment from where global objects should be identified.

Details

This function is *not* meant to be called directly. Instead, the typical usages are:

```
# Evaluate futures via a single background R process on the local machine
plan(cluster, workers = 1)

# Evaluate futures via two background R processes on the local machine
plan(cluster, workers = 2)

# Evaluate futures via a single R process on another machine on on the
# local area network (LAN)
plan(cluster, workers = "raspberrypi")

# Evaluate futures via a single R process running on a remote machine
plan(cluster, workers = "pi.example.org")

# Evaluate futures via four R processes, one running on the local machine,
# two running on LAN machine 'n1' and one on a remote machine
plan(cluster, workers = c("localhost", "n1", "n1", "pi.example.org"))
```

Value

A [ClusterFuture](#).

Examples

```
## Use cluster futures
cl <- parallel::makeCluster(2, timeout = 60)
plan(cluster, workers = cl)

## A global variable
a <- 0

## Create future (explicitly)
f <- future({
  b <- 3
  c <- 2
  a * b * c
})
```

```
## A cluster future is evaluated in a separate process.
## Regardless, changing the value of a global variable will
## not affect the result of the future.
a <- 7
print(a)

v <- value(f)
print(v)
stopifnot(v == 0)

## CLEANUP
parallel::stopCluster(cl)
```

clusterExportSticky *Export globals to the sticky-globals environment of the cluster nodes*

Description

Export globals to the sticky-globals environment of the cluster nodes

Usage

```
clusterExportSticky(cl, globals)
```

Arguments

`cl` (cluster) A cluster object as returned by `parallel::makeCluster()`.
`globals` (list) A named list of sticky globals to be exported.

Details

This requires that the **future** package is installed on the cluster nodes.

Value

(invisible; cluster) The cluster object.

`future`*Create a future*

Description

Creates a future that evaluates an R expression or a future that calls an R function with a set of arguments. How, when, and where these futures are evaluated can be configured using `plan()` such that it is evaluated in parallel on, for instance, the current machine, on a remote machine, or via a job queue on a compute cluster. Importantly, any R code using futures remains the same regardless on these settings and there is no need to modify the code when switching from, say, sequential to parallel processing.

Usage

```
future(  
  expr,  
  envir = parent.frame(),  
  substitute = TRUE,  
  lazy = FALSE,  
  seed = FALSE,  
  globals = TRUE,  
  packages = NULL,  
  stdout = TRUE,  
  conditions = "condition",  
  earlySignal = FALSE,  
  label = NULL,  
  gc = FALSE,  
  ...  
)  
  
futureAssign(  
  x,  
  value,  
  envir = parent.frame(),  
  substitute = TRUE,  
  lazy = FALSE,  
  seed = FALSE,  
  globals = TRUE,  
  packages = NULL,  
  stdout = TRUE,  
  conditions = "condition",  
  earlySignal = FALSE,  
  label = NULL,  
  gc = FALSE,  
  ...,  
  assign.env = envir  
)
```

```

x %<-% value

futureCall(
  FUN,
  args = list(),
  envir = parent.frame(),
  lazy = FALSE,
  seed = FALSE,
  globals = TRUE,
  packages = NULL,
  stdout = TRUE,
  conditions = "condition",
  earlySignal = FALSE,
  label = NULL,
  gc = FALSE,
  ...
)

```

Arguments

expr, value	An R expression .
envir	The environment from where global objects should be identified.
substitute	If TRUE, argument expr is substitute() :ed, otherwise not.
lazy	If FALSE (default), the future is resolved eagerly (starting immediately), otherwise not.
seed	(optional) If TRUE, the random seed, that is, the state of the random number generator (RNG) will be set such that statistically sound random numbers are produced (also during parallelization). If FALSE (default), it is assumed that the future expression does neither need nor use random numbers generation. To use a fixed random seed, specify a L'Ecuyer-CMRG seed (seven integer) or a regular RNG seed (a single integer). If the latter, then a L'Ecuyer-CMRG seed will be automatically created based on the given seed. Furthermore, if FALSE, then the future will be monitored to make sure it does not use random numbers. If it does and depending on the value of option future.rng.onMisuse , the check is ignored, an informative warning, or error will be produced. If seed is NULL, then the effect is as with seed = FALSE but without the RNG check being performed.
globals	(optional) a logical, a character vector, or a named list to control how globals are handled. For details, see section 'Globals used by future expressions' in the help for future() .
packages	(optional) a character vector specifying packages to be attached in the R environment evaluating the future.
stdout	If TRUE (default), then the standard output is captured, and re-outputted when value() is called. If FALSE, any output is silenced (by sinking it to the null device as it is outputted). Using <code>stdout = structure(TRUE, drop = TRUE)</code> causes the captured standard output to be dropped from the future object as soon as it

has been relayed. This can help decrease the overall memory consumed by captured output across futures. Using `stdout = NA` (not recommended) avoids intercepting the standard output; behavior of such unhandled standard output depends on the future

conditions	A character string of conditions classes to be captured and relayed. The default is to relay all conditions, including messages and warnings. To drop all conditions, use <code>conditions = character(0)</code> . Errors are always relayed. Attribute <code>exclude</code> can be used to ignore specific classes, e.g. <code>conditions = structure("condition", exclude = "message")</code> will capture all condition classes except those that inherits from the message class. Using <code>conditions = structure(..., drop = TRUE)</code> causes any captured conditions to be dropped from the future object as soon as it has been relayed, e.g. by <code>value(f)</code> . This can help decrease the overall memory consumed by captured conditions across futures. Using <code>conditions = NULL</code> (not recommended) avoids intercepting conditions, except from errors; behavior of such unhandled conditions depends on the future backend and the environment from which R runs.
earlySignal	Specified whether conditions should be signaled as soon as possible or not.
label	An optional character string label attached to the future.
gc	If TRUE, the garbage collector run (in the process that evaluated the future) only after the value of the future is collected. Exactly when the values are collected may depend on various factors such as number of free workers and whether <code>earlySignal</code> is TRUE (more frequently) or FALSE (less frequently). <i>Some types of futures ignore this argument.</i>
...	Additional arguments passed to <code>Future()</code> .
x	the name of a future variable, which will hold the value of the future expression (as a promise).
assign.env	The environment to which the variable should be assigned.
FUN	A function to be evaluated.
args	A list of arguments passed to function FUN.

Details

The state of a future is either unresolved or resolved. The value of a future can be retrieved using `v <- value(f)`. Querying the value of a non-resolved future will *block* the call until the future is resolved. It is possible to check whether a future is resolved or not without blocking by using `resolved(f)`.

For a future created via a future assignment (`x %<-% value` or `futureAssign("x", value)`), the value is bound to a promise, which when queried will internally call `value()` on the future and which will then be resolved into a regular variable bound to that value. For example, with future assignment `x %<-% value`, the first time variable `x` is queried the call blocks if (and only if) the future is not yet resolved. As soon as it is resolved, and any succeeding queries, querying `x` will immediately give the value.

The future assignment construct `x %<-% value` is not a formal assignment per se, but a binary infix operator on objects `x` and expression `value`. However, by using non-standard evaluation, this constructs can emulate an assignment operator similar to `x <- value`. Due to R's precedence rules of operators, future expressions often need to be explicitly bracketed, e.g. `x %<-% { a + b }`.

The `futureCall()` function works analogously to `do.call()`, which calls a function with a set of arguments. The difference is that `do.call()` returns the value of the call whereas `futureCall()` returns a future.

Value

`f <- future(expr)` creates a **Future** `f` that evaluates expression `expr`, the value of the future is retrieved using `v <- value(f)`.

`x %<-% value` (a future assignment) and `futureAssign("x", value)` create a **Future** that evaluates expression `expr` and binds its value (as a **promise**) to a variable `x`. The value of the future is automatically retrieved when the assigned variable (promise) is queried. The future itself is returned invisibly, e.g. `f <- futureAssign("x", expr)` and `f <- (x %<-% expr)`. Alternatively, the future of a future variable `x` can be retrieved without blocking using `f <- futureOf(x)`. Both the future and the variable (promise) are assigned to environment `assign.env` where the name of the future is `.future_<name>`.

`f <- futureCall(FUN, args)` creates a **Future** `f` that calls function `FUN` with arguments `args`, where the value of the future is retrieved using `x <- value(f)`.

Eager or lazy evaluation

By default, a future is resolved using *eager* evaluation (`lazy = FALSE`). This means that the expression starts to be evaluated as soon as the future is created.

As an alternative, the future can be resolved using *lazy* evaluation (`lazy = TRUE`). This means that the expression will only be evaluated when the value of the future is requested. *Note that this means that the expression may not be evaluated at all - it is guaranteed to be evaluated if the value is requested.*

For future assignments, lazy evaluation can be controlled via the `%lazy%` operator, e.g. `x %<-% { expr } %lazy% TRUE`.

Globals used by future expressions

Global objects (short *globals*) are objects (e.g. variables and functions) that are needed in order for the future expression to be evaluated while not being local objects that are defined by the future expression. For example, in

```
a <- 42
f <- future({ b <- 2; a * b })
```

variable `a` is a global of future assignment `f` whereas `b` is a local variable. In order for the future to be resolved successfully (and correctly), all globals need to be gathered when the future is created such that they are available whenever and wherever the future is resolved.

The default behavior (`globals = TRUE`), is that globals are automatically identified and gathered. More precisely, globals are identified via code inspection of the future expression `expr` and their values are retrieved with environment `envir` as the starting point (basically via `get(global, envir = envir, inherits = TRUE)`). *In most cases, such automatic collection of globals is sufficient and less tedious and error prone than if they are manually specified.*

However, for full control, it is also possible to explicitly specify exactly which the globals are by providing their names as a character vector. In the above example, we could use

```
a <- 42
f <- future({ b <- 2; a * b }, globals = "a")
```

Yet another alternative is to explicitly specify also their values using a named list as in

```
a <- 42
f <- future({ b <- 2; a * b }, globals = list(a = a))
```

or

```
f <- future({ b <- 2; a * b }, globals = list(a = 42))
```

Specifying globals explicitly avoids the overhead added from automatically identifying the globals and gathering their values. Furthermore, if we know that the future expression does not make use of any global variables, we can disable the automatic search for globals by using

```
f <- future({ a <- 42; b <- 2; a * b }, globals = FALSE)
```

Future expressions often make use of functions from one or more packages. As long as these functions are part of the set of globals, the future package will make sure that those packages are attached when the future is resolved. Because there is no need for such globals to be frozen or exported, the future package will not export them, which reduces the amount of transferred objects. For example, in

```
x <- rnorm(1000)
f <- future({ median(x) })
```

variable `x` and `median()` are globals, but only `x` is exported whereas `median()`, which is part of the **stats** package, is not exported. Instead it is made sure that the **stats** package is on the search path when the future expression is evaluated. Effectively, the above becomes

```
x <- rnorm(1000)
f <- future({
  library("stats")
  median(x)
})
```

To manually specify this, one can either do

```
x <- rnorm(1000)
f <- future({
  median(x)
}, globals = list(x = x, median = stats::median)
```

or

```
x <- rnorm(1000)
f <- future({
  library("stats")
  median(x)
}, globals = list(x = x))
```

Both are effectively the same.

Although rarely needed, a combination of automatic identification and manual specification of globals is supported via attributes `add` (to add false negatives) and `ignore` (to ignore false positives) on value `TRUE`. For example, with `globals = structure(TRUE, ignore = "b", add = "a")` any globals automatically identified except `b` will be used in addition to global `a`.

When using future assignments, globals can be specified analogously using the `%globals%` operator, e.g.

```
x <- rnorm(1000)
y %<-% { median(x) } %globals% list(x = x, median = stats::median)
```

Author(s)

The future logo was designed by Dan LaBar and tweaked by Henrik Bengtsson.

See Also

How, when and where futures are resolved is given by the *future strategy*, which can be set by the end user using the `plan()` function. The future strategy must not be set by the developer, e.g. it must not be called within a package.

Examples

```
## Evaluate futures in parallel
plan(multisession)

## Data
x <- rnorm(100)
y <- 2 * x + 0.2 + rnorm(100)
w <- 1 + x ^ 2

## EXAMPLE: Regular assignments (evaluated sequentially)
fitA <- lm(y ~ x, weights = w)      ## with offset
fitB <- lm(y ~ x - 1, weights = w) ## without offset
fitC <- {
  w <- 1 + abs(x) ## Different weights
  lm(y ~ x, weights = w)
}
print(fitA)
print(fitB)
print(fitC)

## EXAMPLE: Future assignments (evaluated in parallel)
fitA %<-% lm(y ~ x, weights = w)    ## with offset
fitB %<-% lm(y ~ x - 1, weights = w) ## without offset
fitC %<-% {
  w <- 1 + abs(x)
  lm(y ~ x, weights = w)
}
```

```

print(fitA)
print(fitB)
print(fitC)

## EXAMPLE: Explicitly create futures (evaluated in parallel)
## and retrieve their values
fA <- future( lm(y ~ x, weights = w) )
fB <- future( lm(y ~ x - 1, weights = w) )
fC <- future({
  w <- 1 + abs(x)
  lm(y ~ x, weights = w)
})
fitA <- value(fA)
fitB <- value(fB)
fitC <- value(fC)
print(fitA)
print(fitB)
print(fitC)

## EXAMPLE: futureCall() and do.call()
x <- 1:100
y0 <- do.call(sum, args = list(x))
print(y0)

f1 <- futureCall(sum, args = list(x))
y1 <- value(f1)
print(y1)

```

future.options

Options used for futures

Description

Below are the R options and environment variables that are used by the **future** package and packages enhancing it.

WARNING: Note that the names and the default values of these options may change in future versions of the package. Please use with care until further notice.

Settings moved to the 'parallelly' package

Several functions have been moved to the **parallelly** package:

- `parallelly::availableCores()`
- `parallelly::availableWorkers()`
- `parallelly::makeClusterMPI()`
- `parallelly::makeClusterPSOCK()`

- `parallely::makeNodePSOCK()`
- `parallely::supportsMulticore()`

The options and environment variables controlling those have been adjusted accordingly to have different prefixes. For example, option `future.fork.enable` has been renamed to `parallely.fork.enable` and the corresponding environment variable `R_FUTURE_FORK_ENABLE` has been renamed to `R_PARALLELY_FORK_ENABLE`. For backward compatibility reasons, the **parallely** package will support both versions for a long foreseeable time. See the [parallely::parallely.options](#) page for the settings.

Options for controlling futures

- ‘`future.plan`’: (character string or future function) Default future strategy plan used unless otherwise specified via `plan()`. This will also be the future plan set when calling `plan("default")`. If not specified, this option may be set when the **future** package is *loaded* if command-line option `--parallel=ncores` (short `-p ncores`) is specified; if `ncores > 1`, then option ‘`future.plan`’ is set to `multisession` otherwise `sequential` (in addition to option ‘`mc.cores`’ being set to `ncores`, if `ncores >= 1`). (Default: `sequential`)
- ‘`future.globals.maxSize`’: (numeric) Maximum allowed total size (in bytes) of global variables identified. Used to prevent too large exports. If set of `+Inf`, then the check for large globals is skipped. (Default: $500 * 1024 ^ 2 = 500$ MiB)
- ‘`future.globals.onReference`’: (**beta feature - may change**) (character string) Controls whether the identified globals should be scanned for so called *references* (e.g. external pointers and connections) or not. It is unlikely that another R process ("worker") can use a global that uses an internal reference of the master R process - we call such objects *non-exportable globals*. If this option is `"error"`, an informative error message is produced if a non-exportable global is detected. If `"warning"`, a warning is produced, but the processing will continue; it is likely that the future will be resolved with a run-time error unless processed in the master R process (e.g. `plan(sequential)` and `plan(multicore)`). If `"ignore"`, no scan is performed. (Default: `"ignore"` but may change)
- ‘`future.resolve.recursive`’: (integer) An integer specifying the maximum recursive depth to which futures should be resolved. If negative, nothing is resolved. If `0`, only the future itself is resolved. If `1`, the future and any of its elements that are futures are resolved, and so on. If `+Inf`, infinite search depth is used. (Default: `0`)
- ‘`future.rng.onMisuse`’: (**beta feature - may change**) (character string) If random numbers are used in futures, then parallel (L’Ecuyer-CMRG) RNG should be used in order to get statistical sound RNGs. The defaults in the future framework assume that *no* random number generation (RNG) is taken place in the future expression because L’Ecuyer-CMRG RNGs come with an unnecessary overhead if not needed. To protect against mistakes, the future framework attempts to detect when random numbers are used despite L’Ecuyer-CMRG RNGs are not in place. If this is detected, and `future.rng.onMisuse = "error"`, then an informative error message is produced. If `"warning"`, then a warning message is produced. If `"ignore"`, no check is performed. (Default: `"warning"`)
- ‘`future.globalenv.onMisuse`’: (**beta feature - may change**) (character string) Assigning variables to the global environment for the purpose of using the variable at a later time makes no sense with futures, because the next future may be evaluated in different R process. To protect against mistakes, the future framework attempts to detect when variables are added to the global environment. If this is detected, and `future.globalenv.onMisuse = "error"`, then

an informative error message is produced. If "warning", then a warning message is produced. If "ignore", no check is performed. (Default: "ignore")

- ‘future.onFutureCondition.keepFuture’: (logical) If TRUE, a FutureCondition keeps a copy of the Future object that triggered the condition. If FALSE, it is dropped. (Default: TRUE)
- ‘future.wait.timeout’: (numeric) Maximum waiting time (in seconds) for a free worker before a timeout error is generated. (Default: 30 * 24 * 60 * 60 (= 30 days))
- ‘future.wait.interval’: (numeric) Initial interval (in seconds) between polls. This controls the polling frequency for finding an available worker when all workers are currently busy. It also controls the polling frequency of resolve(). (Default: 0.01 = 0.01 seconds)
- ‘future.wait.alpha’: (numeric) Positive scale factor used to increase the interval after each poll. (Default: 1.01)

Options for debugging futures

- ‘future.debug’: (logical) If TRUE, extensive debug messages are generated. (Default: FALSE)

Options for controlling package startup

- ‘future.startup.script’: (character vector or a logical) Specifies zero or more future startup scripts to be sourced when the **future** package is *attached*. It is only the first existing script that is sourced. If none of the specified files exist, nothing is sourced - there will be neither a warning nor an error. If this option is not specified, environment variable R_FUTURE_STARTUP_SCRIPT is considered, where multiple scripts may be separated by either a colon (:) or a semicolon (;). If neither is set, or either is set to TRUE, the default is to look for a ‘.future.R’ script in the current directory and then in the user’s home directory. To disable future startup scripts, set the option or the environment variable to FALSE. *Importantly*, this option is *always* set to FALSE if the **future** package is loaded as part of a future expression being evaluated, e.g. in a background process. In other words, they are sourced in the main R process but not in future processes. (Default: TRUE in main R process and FALSE in future processes / during future evaluation)
- ‘future.cmdargs’: (character vector) Overrides `commandArgs()` when the **future** package is *loaded*.

Options for configuring low-level system behaviors

- ‘future.fork.multithreading.enable’ (**beta feature - may change**): (logical) Enable or disable *multi-threading* while using *forked* parallel processing. If FALSE, different multi-thread library settings are overridden such that they run in single-thread mode. Specifically, multi-threading will be disabled for OpenMP (which requires the **RhpcBLASctl** package) and for **RcppParallel**. If TRUE, or not set (the default), multi-threading is allowed. Parallelization via multi-threaded processing (done in native code by some packages and external libraries) while at the same time using forked (aka "multicore") parallel processing is known to unstable. Note that this is not only true when using `plan(multicore)` but also when using, for instance, `mclapply()` of the **parallel** package. (Default: not set)
- ‘future.output.windows.reencode’: (logical) Enable or disable re-encoding of UTF-8 symbols that were incorrectly encoded while captured. In R (< 4.2.0) and on older versions of MS Windows, R cannot capture UTF-8 symbols as-is when they are captured from the standard output. For examples, a UTF-8 check mark symbol (“\u2713”) would be relayed as "<U+2713>" (a

string with eight ASCII characters). Setting this option to TRUE will cause `value()` to attempt to recover the intended UTF-8 symbols from `<U+nnnn>` string components, if, and only if, the string was captured by a future resolved on MS Windows. (Default: TRUE)

See also [parallexly::parallexly.options](#).

Options for demos

`'future.demo.mandelbrot.region'`: (integer) Either a named list of `mandelbrot()` arguments or an integer in {1, 2, 3} specifying a predefined Mandelbrot region. (Default: 1L)

`'future.demo.mandelbrot.nrow'`: (integer) Number of rows and columns of tiles. (Default: 3L)

Deprecated or for internal prototyping

The following options exist only for troubleshooting purposes and must not be used in production. If used, there is a risk that the results are non-reproducible if processed elsewhere. To lower the risk of them being used by mistake, they are marked as deprecated and will produce warnings if set.

`'future.globals.onMissing'`: (character string) Action to take when non-existing global variables ("globals" or "unknowns") are identified when the future is created. If "error", an error is generated immediately. If "ignore", no action is taken and an attempt to evaluate the future expression will be made. The latter is useful when there is a risk for false-positive globals being identified, e.g. when future expression contains non-standard evaluation (NSE). (Default: "ignore")

`'future.globals.method'`: (character string) Method used to identify globals. For details, see `globalsOf()`. (Default: "ordered")

`'future.globals.resolve'`: (logical) If TRUE, globals that are `Future` objects (typically created as *explicit* futures) will be resolved and have their values (using `value()`) collected. Because searching for unresolved futures among globals (including their content) can be expensive, the default is not to do it and instead leave it to the run-time checks that assert proper ownership when resolving futures and collecting their values. (Default: FALSE)

Environment variables that set R options

All of the above R `'future.*'` options can be set by corresponding environment variable `R_FUTURE_*` when the **future** package is loaded. This means that those environment variables must be set before the **future** package is loaded in order to have an effect. For example, if `R_FUTURE_RNG_ONMISUSE = "ignore"`, then option `'future.rng.onMisuse'` is set to "ignore" (character string). Similarly, if `R_FUTURE_GLOBALS_MAXSIZE = "50000000"`, then option `'future.globals.maxSize'` is set to 50000000 (numeric).

See Also

To set R options or environment variables when R starts (even before the **future** package is loaded), see the [Startup](#) help page. The **startup** package provides a friendly mechanism for configuring R's startup process.

Examples

```
# Allow at most 5 MB globals per futures
options(future.globals.maxSize = 5e6)

# Be strict; catch all RNG mistakes
options(future.rng.onMisuse = "error")
```

futureOf	<i>Get the future of a future variable</i>
----------	--

Description

Get the future of a future variable that has been created directly or indirectly via `future()`.

Usage

```
futureOf(
  var = NULL,
  envir = parent.frame(),
  mustExist = TRUE,
  default = NA,
  drop = FALSE
)
```

Arguments

var	the variable. If NULL, all futures in the environment are returned.
envir	the environment where to search from.
mustExist	If TRUE and the variable does not exist, then an informative error is thrown, otherwise NA is returned.
default	the default value if future was not found.
drop	if TRUE and var is NULL, then returned list only contains futures, otherwise also default values.

Value

A `Future` (or default). If var is NULL, then a named list of `Future`s are returned.

Examples

```
a %<-% { 1 }

f <- futureOf(a)
print(f)
```



```
b %<-% { 2 }

f <- futureOf(b)
print(f)

## All futures
fs <- futureOf()
print(fs)

## Futures part of environment
env <- new.env()
env$c %<-% { 3 }

f <- futureOf(env$c)
print(f)

f2 <- futureOf(c, envir = env)
print(f2)

f3 <- futureOf("c", envir = env)
print(f3)

fs <- futureOf(envir = env)
print(fs)
```

futures

Get all futures in a container

Description

Gets all futures in an environment, a list, or a list environment and returns an object of the same class (and dimensions). Non-future elements are returned as is.

Usage

```
futures(x, ...)
```

Arguments

x	An environment, a list, or a list environment.
...	Not used.

Details

This function is useful for retrieve futures that were created via future assignments (%<-%) and therefore stored as promises. This function turns such promises into standard Future objects.

Value

An object of same type as `x` and with the same names and/or dimensions, if set.

<code>futureSessionInfo</code>	<i>Get future-specific session information and validate current backend</i>
--------------------------------	---

Description

Get future-specific session information and validate current backend

Usage

```
futureSessionInfo(test = TRUE, anonymize = TRUE)
```

Arguments

<code>test</code>	If TRUE, one or more futures are created to query workers and validate their information.
<code>anonymize</code>	If TRUE, user names and host names are anonymized.

Value

Nothing.

Examples

```
plan(multisession, workers = 2)
futureSessionInfo()
plan(sequential)
```

<code>multicore</code>	<i>Create a multicore future whose value will be resolved asynchronously in a forked parallel process</i>
------------------------	---

Description

A multicore future is a future that uses multicore evaluation, which means that its *value is computed and resolved in parallel in another process*.

Usage

```
multicore(
  ...,
  workers = availableCores(constraints = "multicore"),
  envir = parent.frame()
)
```

Arguments

...	Additional arguments passed to <code>Future()</code> .
workers	The number of parallel processes to use. If a function, it is called without arguments <i>when the future is created</i> and its value is used to configure the workers.
envir	The environment from where global objects should be identified.

Details

This function is *not* meant to be called directly. Instead, the typical usages are:

```
# Evaluate futures in parallel on the local machine via as many forked
# processes as available to the current R process
plan(multicore)

# Evaluate futures in parallel on the local machine via two forked processes
plan(multicore, workers = 2)
```

Value

A **MulticoreFuture**. If `workers == 1`, then all processing using done in the current/main R session and we therefore fall back to using a sequential future. To override this fallback, use `workers = I(1)`. This is also the case whenever multicore processing is not supported, e.g. on Windows.

Support for forked ("multicore") processing

Not all operating systems support process forking and thereby not multicore futures. For instance, forking is not supported on Microsoft Windows. Moreover, process forking may break some R environments such as RStudio. Because of this, the future package disables process forking also in such cases. See `parallely::supportsMulticore()` for details. Trying to create multicore futures on non-supported systems or when forking is disabled will result in multicore futures falling back to becoming **sequential** futures. If used in RStudio, there will be an informative warning:

```
> plan(multicore)
Warning message:
In supportsMulticoreAndRStudio(...) :
  [ONE-TIME WARNING] Forked processing ('multicore') is not supported when
  running R from RStudio because it is considered unstable. For more details,
  how to control forked processing or not, and how to silence this warning in
  future R sessions, see ?parallely::supportsMulticore
```

See Also

For processing in multiple background R sessions, see **multisession** futures.

Use `parallely::availableCores()` to see the total number of cores that are available for the current R session. Use `availableCores("multicore") > 1L` to check whether multicore futures are supported or not on the current system.

Examples

```

## Use multicore futures
plan(multicore)

## A global variable
a <- 0

## Create future (explicitly)
f <- future({
  b <- 3
  c <- 2
  a * b * c
})

## A multicore future is evaluated in a separate forked
## process. Changing the value of a global variable
## will not affect the result of the future.
a <- 7
print(a)

v <- value(f)
print(v)
stopifnot(v == 0)

```

multisession	<i>Create a multisession future whose value will be resolved asynchronously in a parallel R session</i>
--------------	---

Description

A multisession future is a future that uses multisession evaluation, which means that its *value is computed and resolved in parallel in another R session*.

Usage

```

multisession(
  ...,
  workers = availableCores(),
  lazy = FALSE,
  rscript_libs = .libPaths(),
  envir = parent.frame()
)

```

Arguments

...	Additional arguments passed to Future() .
workers	The number of parallel processes to use. If a function, it is called without arguments <i>when the future is created</i> and its value is used to configure the workers.

lazy	If FALSE (default), the future is resolved eagerly (starting immediately), otherwise not.
rscript_libs	A character vector of R package library folders that the workers should use. The default is <code>.libPaths()</code> so that multisession workers inherits the same library path as the main R session. To avoid this, use <code>plan(multisession, ..., rscript_libs = NULL)</code> . <i>Important: Note that the library path is set on the workers when they are created, i.e. when <code>plan(multisession)</code> is called. Any changes to <code>.libPaths()</code> in the main R session after the workers have been created will have no effect.</i> This is passed down as-is to <code>parallely::makeClusterPSOCK()</code> .
envir	The environment from where global objects should be identified.

Details

This function is *not* meant to be called directly. Instead, the typical usages are:

```
# Evaluate futures in parallel on the local machine via as many background
# processes as available to the current R process
plan(multisession)
```

```
# Evaluate futures in parallel on the local machine via two background
# processes
plan(multisession, workers = 2)
```

The background R sessions (the "workers") are created using `makeClusterPSOCK()`.

For the total number of R sessions available including the current/main R process, see `parallely::availableCores()`.

A multisession future is a special type of cluster future.

Value

A [MultisessionFuture](#). If `workers == 1`, then all processing using done in the current/main R session and we therefore fall back to using a lazy future. To override this fallback, use `workers = I(1)`.

See Also

For processing in multiple forked R sessions, see [multicore](#) futures.

Use `parallely::availableCores()` to see the total number of cores that are available for the current R session.

Examples

```
## Use multisession futures
plan(multisession)

## A global variable
a <- 0
```

```

## Create future (explicitly)
f <- future({
  b <- 3
  c <- 2
  a * b * c
})

## A multisession future is evaluated in a separate R session.
## Changing the value of a global variable will not affect
## the result of the future.
a <- 7
print(a)

v <- value(f)
print(v)
stopifnot(v == 0)

## Explicitly close multisession workers by switching plan
plan(sequential)

```

nbrOfWorkers

Get the number of workers available

Description

Get the number of workers available

Usage

```
nbrOfWorkers(evaluator = NULL)
```

```
nbrOfFreeWorkers(evaluator = NULL, background = FALSE, ...)
```

Arguments

evaluator	A future evaluator function. If NULL (default), the current evaluator as returned by <code>plan()</code> is used.
background	If TRUE, only workers that can process a future in the background are considered. If FALSE, also workers running in the main R process are considered, e.g. when using the 'sequential' backend.
...	Not used; reserved for future use.

Value

`nbrOfWorkers()` returns a positive number in 1, 2, 3, ..., which for some future backends may also be +Inf.

`nbrOfFreeWorkers()` returns a non-negative number in 0, 1, 2, 3, ... which is less than or equal to `nbrOfWorkers()`.

Examples

```

plan(multisession)
nbrOfWorkers() ## == availableCores()

plan(sequential)
nbrOfWorkers() ## == 1

```

plan

Plan how to resolve a future

Description

This function allows *the user* to plan the future, more specifically, it specifies how `future():s` are resolved, e.g. sequentially or in parallel.

Usage

```

plan(
  strategy = NULL,
  ...,
  substitute = TRUE,
  .skip = FALSE,
  .call = TRUE,
  .cleanup = TRUE,
  .init = TRUE
)

```

Arguments

<code>strategy</code>	The evaluation function (or name of it) to use for resolving a future. If <code>NULL</code> , then the current strategy is returned.
<code>...</code>	Additional arguments overriding the default arguments of the evaluation function. Which additional arguments are supported depends on what evaluation function is used, e.g. several support argument <code>workers</code> but not all. For details, see the individual functions of which some are linked to below.
<code>substitute</code>	If <code>TRUE</code> , the strategy expression is <code>substitute():d</code> , otherwise not.
<code>.skip</code>	(internal) If <code>TRUE</code> , then attempts to set a strategy that is the same as what is currently in use, will be skipped.
<code>.call</code>	(internal) Used for recording the call to this function.
<code>.cleanup</code>	(internal) Used to stop implicitly started clusters.
<code>.init</code>	(internal) Used to initiate workers.

Details

The default strategy is `sequential`, but the default can be configured by option `'future.plan'` and, if that is not set, system environment variable `R_FUTURE_PLAN`. To reset the strategy back to the default, use `plan("default")`.

Value

If a new strategy is chosen, then the previous one is returned (invisible), otherwise the current one is returned (visibly).

Built-in evaluation strategies

The **future** package provides the following built-in backends:

- **sequential**: Resolves futures sequentially in the current R process, e.g. `plan(sequential)`.
- **multisession**: Resolves futures asynchronously (in parallel) in separate R sessions running in the background on the same machine, e.g. `plan(multisession)` and `plan(multisession, workers = 2)`.
- **multicore**: Resolves futures asynchronously (in parallel) in separate *forked* R processes running in the background on the same machine, e.g. `plan(multicore)` and `plan(multicore, workers = 2)`. This backend is not supported on Windows.
- **cluster**: Resolves futures asynchronously (in parallel) in separate R sessions running typically on one or more machines, e.g. `plan(cluster)`, `plan(cluster, workers = 2)`, and `plan(cluster, workers = c("n1", "n1", "n2", "server.remote.org"))`.

Other packages provide additional evaluation strategies. For example, the **future.callr** package implements an alternative to the `multisession` backend on top of the **callr** package, e.g. `plan(future.callr::callr, workers = 2)`. Another example is the **future.batchtools** package, which implements, on top of the **batchtools** package, e.g. `plan(future.batchtools::batchtools_slurm)`. These types of futures are resolved via job schedulers, which typically are available on high-performance compute (HPC) clusters, e.g. LSF, Slurm, TORQUE/PBS, Sun Grid Engine, and OpenLava.

To "close" any background workers (e.g. `multisession`), change the plan to something different; `plan(sequential)` is recommended for this.

For package developers

Please refrain from modifying the future strategy inside your packages / functions, i.e. do not call `plan()` in your code. Instead, leave the control on what backend to use to the end user. This idea is part of the core philosophy of the future framework - as a developer you can never know what future backends the user has access to. Moreover, by not making any assumptions about what backends are available, your code will also work automatically with any new backends developed after you wrote your code.

If you think it is necessary to modify the future strategy within a function, then make sure to undo the changes when exiting the function. This can be done using:

```
oplan <- plan(new_set_of_strategies)
on.exit(plan(oplan), add = TRUE)
[...]
```

This is important because the end-user might have already set the future strategy elsewhere for other purposes and will most likely not know that calling your function will break their setup. *Remember, your package and its functions might be used in a greater context where multiple packages and functions are involved and those might also rely on the future framework, so it is important to avoid stepping on others' toes.*

Using plan() in scripts and vignettes

When writing scripts or vignettes that uses futures, try to place any call to `plan()` as far up (as early on) in the code as possible. This will help users to quickly identify where the future plan is set up and allow them to modify it to their computational resources. Even better is to leave it to the user to set the `plan()` prior to `source()`:ing the script or running the vignette. If a `‘.future.R’` exists in the current directory and / or in the user’s home directory, it is sourced when the **future** package is *loaded*. Because of this, the `‘.future.R’` file provides a convenient place for users to set the `plan()`. This behavior can be controlled via an R option - see [future options](#) for more details.

Examples

```
a <- b <- c <- NA_real_

# An sequential future
plan(sequential)
f <- future({
  a <- 7
  b <- 3
  c <- 2
  a * b * c
})
y <- value(f)
print(y)
str(list(a = a, b = b, c = c)) ## All NAs

# A sequential future with lazy evaluation
plan(sequential)
f <- future({
  a <- 7
  b <- 3
  c <- 2
  a * b * c
}, lazy = TRUE)
y <- value(f)
print(y)
str(list(a = a, b = b, c = c)) ## All NAs

# A multicore future (specified as a string)
plan("multicore")
f <- future({
  a <- 7
  b <- 3
  c <- 2
  a * b * c
})
y <- value(f)
print(y)
str(list(a = a, b = b, c = c)) ## All NAs
```

```
## Multisession futures gives an error on R CMD check on
## Windows (but not Linux or macOS) for unknown reasons.
## The same code works in package tests.

# A multisession future (specified via a string variable)
plan("future::multisession")
f <- future({
  a <- 7
  b <- 3
  c <- 2
  a * b * c
})
y <- value(f)
print(y)
str(list(a = a, b = b, c = c)) ## All NAs

## Explicitly specifying number of workers
## (default is parallelly::availableCores())
plan(multicore, workers = 2)
message("Number of parallel workers: ", nbrOfWorkers())

## Explicitly close multisession workers by switching plan
plan(sequential)
```

re-exports

Functions Moved to 'parallelly'

Description

The following function used to be part of **future** but has since been migrated to **parallelly**. The migration started with **future** 1.20.0 (November 2020). They were moved because they are also useful outside of the **future** framework.

Details

- [parallelly::as.cluster\(\)](#)
- [parallelly::autoStopCluster\(\)](#)
- [parallelly::availableCores\(\)](#)
- [parallelly::availableWorkers\(\)](#)
- [parallelly::makeClusterMPI\(\)](#)
- [parallelly::makeClusterPSOCK\(\)](#)
- [parallelly::makeNodePSOCK\(\)](#)

- `parallely::supportsMulticore()`

For backward-compatible reasons, these functions remain available as exact copies also from this package (as re-exports). For example,

```
cl <- parallely::makeClusterPSOCK(2)
```

can still be accessed as:

```
cl <- future::makeClusterPSOCK(2)
```

resolve

Resolve one or more futures synchronously

Description

This function provides an efficient mechanism for waiting for multiple futures in a container (e.g. list or environment) to be resolved while in the meanwhile retrieving values of already resolved futures.

Usage

```
resolve(
  x,
  idxs = NULL,
  recursive = 0,
  result = FALSE,
  stdout = FALSE,
  signal = FALSE,
  force = FALSE,
  sleep = getOption("future.wait.interval", 0.01),
  ...
)
```

Arguments

<code>x</code>	A Future to be resolved, or a list, an environment, or a list environment of futures to be resolved.
<code>idxs</code>	(optional) integer or logical index specifying the subset of elements to check.
<code>recursive</code>	A non-negative number specifying how deep of a recursion should be done. If TRUE, an infinite recursion is used. If FALSE or zero, no recursion is performed.
<code>result</code>	(internal) If TRUE, the results are <i>retrieved</i> , otherwise not. Note that this only collects the results from the parallel worker, which can help lower the overall latency if there are multiple concurrent futures. This does <i>not</i> return the collected results.

stdout	(internal) If TRUE, captured standard output is relayed, otherwise not.
signal	(internal) If TRUE, captured conditions are relayed, otherwise not.
force	(internal) If TRUE, captured standard output and captured conditions already relayed is relayed again, otherwise not.
sleep	Number of seconds to wait before checking if futures have been resolved since last time.
...	Not used.

Details

This function is resolves synchronously, i.e. it blocks until x and any containing futures are resolved.

Value

Returns x (regardless of subsetting or not). If signal is TRUE and one of the futures produces an error, then that error is produced.

See Also

To resolve a future *variable*, first retrieve its [Future](#) object using `futureOf()`, e.g. `resolve(futureOf(x))`.

 resolved

Check whether a future is resolved or not

Description

Check whether a future is resolved or not

Usage

```
resolved(x, ...)
```

Arguments

x	A Future , a list, or an environment (which also includes list environment).
...	Not used.

Details

This method needs to be implemented by the class that implement the Future API. The implementation should return either TRUE or FALSE and must never throw an error (except for [FutureError:s](#) which indicate significant, often unrecoverable infrastructure problems). It should also be possible to use the method for polling the future until it is resolved (without having to wait infinitely long), e.g. `while (!resolved(future)) Sys.sleep(5)`.

Value

A logical of the same length and dimensions as x. Each element is TRUE unless the corresponding element is a non-resolved future in case it is FALSE.

sequential	Create a sequential future whose value will be in the current R session
------------	---

Description

A sequential future is a future that is evaluated sequentially in the current R session similarly to how R expressions are evaluated in R. The only difference to R itself is that globals are validated by default just as for all other types of futures in this package.

Usage

```
sequential(..., envir = parent.frame())
```

Arguments

... Additional arguments passed to [Future\(\)](#).
envir The [environment](#) from where global objects should be identified.

Details

This function is *not* meant to be called directly. Instead, the typical usages are:

```
# Evaluate futures sequentially in the current R process  
plan(sequential)
```

Value

A [SequentialFuture](#).

Examples

```
## Use sequential futures  
plan(sequential)  
  
## A global variable  
a <- 0  
  
## Create a sequential future  
f <- future({  
  b <- 3  
  c <- 2  
  a * b * c  
})  
  
## Since 'a' is a global variable in future 'f' which  
## is eagerly resolved (default), this global has already  
## been resolved / incorporated, and any changes to 'a'  
## at this point will not affect the value of 'f'.  
a <- 7
```

```
print(a)

v <- value(f)
print(v)
stopifnot(v == 0)
```

tweak	<i>Tweak a future function by adjusting its default arguments</i>
-------	---

Description

Tweak a future function by adjusting its default arguments

Usage

```
tweak(strategy, ..., penvir = parent.frame())
```

Arguments

strategy	An existing future function or the name of one.
...	Named arguments to replace the defaults of existing arguments.
penvir	The environment used when searching for a future function by its name.

Value

a future function.

See Also

Use [plan\(\)](#) to set a future to become the new default strategy.

value	<i>The value of a future or the values of all elements in a container</i>
-------	---

Description

Gets the value of a future or the values of all elements (including futures) in a container such as a list, an environment, or a list environment. If one or more futures is unresolved, then this function blocks until all queried futures are resolved.

Usage

```

value(...)

## S3 method for class 'Future'
value(future, stdout = TRUE, signal = TRUE, ...)

## S3 method for class 'list'
value(x, stdout = TRUE, signal = TRUE, ...)

## S3 method for class 'listenv'
value(x, stdout = TRUE, signal = TRUE, ...)

## S3 method for class 'environment'
value(x, stdout = TRUE, signal = TRUE, ...)

```

Arguments

- ... All arguments used by the S3 methods.
- future, x A [Future](#), an environment, a list, or a list environment.
- stdout If TRUE, standard output captured while resolving futures is relayed, otherwise not.
- signal If TRUE, [conditions](#) captured while resolving futures are relayed, otherwise not.

Value

value() of a Future object returns the value of the future, which can be any type of R object.

value() of a list, an environment, or a list environment returns an object with the same number of elements and of the same class. Names and dimension attributes are preserved, if available. All future elements are replaced by their corresponding value() values. For all other elements, the existing object is kept as-is.

If signal is TRUE and one of the futures produces an error, then that error is produced.

%conditions%	<i>Control whether standard output should be captured or not</i>
--------------	--

Description

Control whether standard output should be captured or not

Usage

```
fassignment %conditions% capture
```

Arguments

- fassignment The future assignment, e.g. x %<-% { expr }.
- capture If TRUE, the standard output will be captured, otherwise not.

`%globals%`*Specify globals and packages for a future assignment*

Description

Specify globals and packages for a future assignment

Usage

```
fassignment %globals% globals
fassignment %packages% packages
```

Arguments

<code>fassignment</code>	The future assignment, e.g. <code>x %<-% { expr }</code> .
<code>globals</code>	(optional) a logical, a character vector, or a named list to control how globals are handled. For details, see section 'Globals used by future expressions' in the help for future() .
<code>packages</code>	(optional) a character vector specifying packages to be attached in the R environment evaluating the future.

`%label%`*Specify label for a future assignment*

Description

Specify label for a future assignment

Usage

```
fassignment %label% label
```

Arguments

<code>fassignment</code>	The future assignment, e.g. <code>x %<-% { expr }</code> .
<code>label</code>	An optional character string label attached to the future.

`%lazy%`

Control lazy / eager evaluation for a future assignment

Description

Control lazy / eager evaluation for a future assignment

Usage

```
fassignment %lazy% lazy
```

Arguments

<code>fassignment</code>	The future assignment, e.g. <code>x %<-% { expr }</code> .
<code>lazy</code>	If FALSE (default), the future is resolved eagerly (starting immediately), otherwise not.

`%plan%`

Use a specific plan for a future assignment

Description

Use a specific plan for a future assignment

Usage

```
fassignment %plan% strategy
```

Arguments

<code>fassignment</code>	The future assignment, e.g. <code>x %<-% { expr }</code> .
<code>strategy</code>	The mechanism for how the future should be resolved. See plan() for further details.

See Also

The [plan\(\)](#) function sets the default plan for all futures.

%seed%	<i>Set random seed for future assignment</i>
--------	--

Description

Set random seed for future assignment

Usage

```
fassignment %seed% seed
```

Arguments

fassignment	The future assignment, e.g. <code>x %<-% { expr }</code> .
seed	(optional) If TRUE, the random seed, that is, the state of the random number generator (RNG) will be set such that statistically sound random numbers are produced (also during parallelization). If FALSE (default), it is assumed that the future expression does neither need nor use random numbers generation. To use a fixed random seed, specify a L'Ecuyer-CMRG seed (seven integer) or a regular RNG seed (a single integer). If the latter, then a L'Ecuyer-CMRG seed will be automatically created based on the given seed. Furthermore, if FALSE, then the future will be monitored to make sure it does not use random numbers. If it does and depending on the value of option <code>future.rng.onMisuse</code> , the check is ignored, an informative warning, or error will be produced. If seed is NULL, then the effect is as with <code>seed = FALSE</code> but without the RNG check being performed.

%stdout%	<i>Control whether standard output should be captured or not</i>
----------	--

Description

Control whether standard output should be captured or not

Usage

```
fassignment %stdout% capture
```

Arguments

fassignment	The future assignment, e.g. <code>x %<-% { expr }</code> .
capture	If TRUE, the standard output will be captured, otherwise not.

`%tweak%` *Temporarily tweaks the arguments of the current strategy*

Description

Temporarily tweaks the arguments of the current strategy

Usage

`fassignment %tweak% tweaks`

Arguments

`fassignment` The future assignment, e.g. `x %<-% { expr }`.

`tweaks` A named list (or vector) with arguments that should be changed relative to the current strategy.

Index

* internals

- clusterExportSticky, 5
- .future.R, 25
- .future.R (future.options), 12
- %->% (future), 6
- %<-% (future), 6
- %packages% (%globals%), 32
- %conditions%, 31
- %globals%, 11, 32
- %label%, 32
- %lazy%, 33
- %plan%, 33
- %seed%, 34
- %stdout%, 34
- %tweak%, 35

- as.cluster (re-exports), 26
- autoStopCluster (re-exports), 26
- availableCores, 19
- availableCores (re-exports), 26
- availableWorkers (re-exports), 26

- backtrace, 2

- cluster, 3, 4, 24
- clusterExportSticky, 5
- ClusterFuture, 4
- ClusterFuture(), 3
- commandArgs, 14
- conditions, 28, 31

- do.call, 9

- environment, 4, 7, 8, 19, 21, 29
- expression, 7

- function, 8
- Future, 9, 15, 16, 27, 28, 31
- future, 6
- future options, 25
- Future(), 8, 19, 20, 29

- future(), 7, 16, 23, 32
- future.cmdargs (future.options), 12
- future.debug (future.options), 12
- future.demo.mandelbrot.nrow (future.options), 12
- future.demo.mandelbrot.region (future.options), 12
- future.fork.multithreading.enable (future.options), 12
- future.globalenv.onMisuse (future.options), 12
- future.globals.maxSize (future.options), 12
- future.globals.method (future.options), 12
- future.globals.onMissing (future.options), 12
- future.globals.onReference (future.options), 12
- future.globals.resolve (future.options), 12
- future.journal (future.options), 12
- future.onFutureCondition.keepFuture (future.options), 12
- future.options, 12
- future.output.windows.reencode (future.options), 12
- future.plan (future.options), 12
- future.resolve.recursive (future.options), 12
- future.rng.onMisuse, 7, 34
- future.rng.onMisuse (future.options), 12
- future.startup.script (future.options), 12
- future.wait.alpha (future.options), 12
- future.wait.interval (future.options), 12
- future.wait.timeout (future.options), 12
- futureAssign (future), 6

- futureCall (future), 6
- FutureError, 28
- futureOf, 9, 16
- futureOf(), 28
- futures, 17
- futureSessionInfo, 18

- globalsOf, 15

- list, 8
- list environment, 28

- makeClusterMPI (re-exports), 26
- makeClusterPSOCK, 4
- makeClusterPSOCK (re-exports), 26
- makeClusterPSOCK(), 21
- makeNodePSOCK (re-exports), 26
- mandelbrot(), 15
- mclapply, 14
- multicore, 18, 21, 24
- MulticoreFuture, 19
- multisession, 19, 20, 24
- MultisessionFuture, 21

- nbrOfFreeWorkers (nbrOfWorkers), 22
- nbrOfWorkers, 22

- parallel::makeCluster(), 5
- parallelly::as.cluster(), 26
- parallelly::autoStopCluster(), 26
- parallelly::availableCores(), 12, 19, 21, 26
- parallelly::availableWorkers(), 12, 26
- parallelly::makeClusterMPI(), 12, 26
- parallelly::makeClusterPSOCK(), 12, 21, 26
- parallelly::makeNodePSOCK(), 13, 26
- parallelly::parallelly.options, 13, 15
- parallelly::supportsMulticore(), 13, 19, 27
- plan, 23
- plan(), 6, 11, 13, 22, 30, 33
- promise, 9

- R_FUTURE_DEBUG (future.options), 12
- R_FUTURE_DEMO_MANDELBROT_NROW (future.options), 12
- R_FUTURE_DEMO_MANDELBROT_REGION (future.options), 12
- R_FUTURE_FORK_MULTITHREADING_ENABLE (future.options), 12
- R_FUTURE_GLOBAENV_ONMISUSE (future.options), 12
- R_FUTURE_GLOBALS_MAXSIZE (future.options), 12
- R_FUTURE_GLOBALS_METHOD (future.options), 12
- R_FUTURE_GLOBALS_ONMISSING (future.options), 12
- R_FUTURE_GLOBALS_ONREFERENCE (future.options), 12
- R_FUTURE_GLOBALS_RESOLVE (future.options), 12
- R_FUTURE_JOURNAL (future.options), 12
- R_FUTURE_ONFUTURECONDITION_KEEPFUTURE (future.options), 12
- R_FUTURE_OUTPUT_WINDOWS_REENCODE (future.options), 12
- R_FUTURE_PLAN (future.options), 12
- R_FUTURE_RESOLVE_RECURSIVE (future.options), 12
- R_FUTURE_RESOLVED_TIMEOUT (future.options), 12
- R_FUTURE_RNG_ONMISUSE (future.options), 12
- R_FUTURE_STARTUP_SCRIPT (future.options), 12
- R_FUTURE_WAIT_ALPHA (future.options), 12
- R_FUTURE_WAIT_INTERVAL (future.options), 12
- R_FUTURE_WAIT_TIMEOUT (future.options), 12
- re-exports, 26
- resolve, 27
- resolved, 8, 28

- sequential, 19, 23, 24, 29
- SequentialFuture, 29
- Startup, 15
- substitute, 7
- supportsMulticore (re-exports), 26

- tweak, 30

- uniprocess (sequential), 29

- value, 8, 30
- value(), 8