

Errors, Logs and Debugging in *BiocParallel*

Valerie Obenchain and Martin Morgan

Edited: December 16, 2015; Compiled: June 25, 2024

Contents

1	Introduction	1
2	Error Handling	2
2.1	Messages and warnings	2
2.2	Catching errors	2
2.3	Identify failures with <code>bpok()</code>	6
2.4	Rerun failed tasks with <code>BPRED0</code>	7
3	Logging	8
3.1	Parameters.	8
3.2	Setting a threshold	9
3.3	Log files	12
4	Worker timeout	13
5	Debugging.	13
5.1	Accessing the traceback	14
5.2	Adding debug messages	15
5.3	Local debugging with <code>SerialParam</code>	18
6	<code>sessionInfo()</code>	19

1 Introduction

This vignette is part of the *BiocParallel* package and focuses on error handling and logging. A section at the end demonstrates how the two can be used together as part of an effective debugging routine.

BiocParallel provides a unified interface to the parallel infrastructure in several packages including *snow*, *parallel*, *batchtools* and *foreach*. When implementing error handling in *BiocParallel* the primary goals were to enable the return of partial results when an error is thrown (vs just the error) and to establish logging on the workers. In cases where error

Errors, Logs and Debugging in *BiocParallel*

handling existed, such as *batchtools* and *foreach*, those behaviors were preserved. Clusters created with *snow* and *parallel* now have flexible error handling and logging available through *SnowParam* and *MulticoreParam* objects.

In this document the term “job” is used to describe a single call to a `bp*apply` function (e.g., the `X` in `bp*apply`). A “job” consists of one or more “tasks”, where each “task” is run separately on a worker.

The *BiocParallel* package is available at bioconductor.org and can be downloaded via `BiocManager::install`:

```
if (!requireNamespace("BiocManager", quietly = TRUE))
  install.packages("BiocManager")
BiocManager::install("BiocParallel")
```

Load the package:

```
library(BiocParallel)
```

2 Error Handling

2.1 Messages and warnings

BiocParallel captures messages and warnings in each job, returning the output to the manager and reporting these to the user after the completion of the entire operation. Thus

```
res <- bplapply(1:2, function(i) { message(i); Sys.sleep(3) })
```

reports messages only after the entire `bplapply()` is complete.

It may be desired to output messages immediately. Do this using `sink()`, as in the following example:

```
res <- bplapply(1:2, function(i) {
  sink(NULL, type = "message")
  message(i)
  Sys.sleep(3)
})
```

This could be confusing when multiple workers write messages at the same time –the messages will be interleaved in an arbitrary way – or when the workers are not all running on the same computer (e.g., with `SnowParam()`) so should not be used in package code.

2.2 Catching errors

By default, *BiocParallel* attempts all computations and returns any warnings and errors along with successful results. The `stop.on.error` field controls if the job is terminated as soon as one task throws an error. This is useful when debugging or when running large jobs (many tasks) and you want to be notified of an error before all runs complete.

`stop.on.error` is `TRUE` by default.

Errors, Logs and Debugging in *BiocParallel*

```
param <- SnowParam()
param

## class: SnowParam
##  bpisup: FALSE; bpnworkers: 2; bptasks: 0; bpjobname: BPJOB
##  bplog: FALSE; bpthreshold: INFO; bpstopOnError: TRUE
##  bpRNGseed: ; bptimeout: NA; bpprogressbar: FALSE
##  bpexportglobals: TRUE; bpexportvariables: TRUE; bpforceGC: FALSE
##  bpfallback: TRUE
##  bplogdir: NA
##  bpresultdir: NA
##  cluster type: SOCK
```

The field can be set when constructing the param or modified with the `bpstopOnError` accessor.

```
param <- SnowParam(2, stop.on.error = TRUE)
param

## class: SnowParam
##  bpisup: FALSE; bpnworkers: 2; bptasks: 0; bpjobname: BPJOB
##  bplog: FALSE; bpthreshold: INFO; bpstopOnError: TRUE
##  bpRNGseed: ; bptimeout: NA; bpprogressbar: FALSE
##  bpexportglobals: TRUE; bpexportvariables: TRUE; bpforceGC: FALSE
##  bpfallback: TRUE
##  bplogdir: NA
##  bpresultdir: NA
##  cluster type: SOCK

bpstopOnError(param) <- FALSE
```

In this example `X` is length 6. By default, the elements of `X` are divided as evenly as possible over the number of workers and run in chunks. The number of tasks is set equal to the length of `X` which forces each element of `X` to be executed separately (6 tasks).

```
X <- list(1, "2", 3, 4, 5, 6)
param <- SnowParam(3, tasks = length(X), stop.on.error = TRUE)
```

Tasks 1, 2, and 3 are assigned to the three workers, and are evaluated. Task 2 fails, stopping further computation. All successfully completed tasks are returned and can be accessed by 'bpresult'. Usually, this means that the results of tasks 1, 2, and 3 will be returned.

```
result <- tryCatch({
  bplapply(X, sqrt, BPPARAM = param)
}, error=identity)
result

## <bplist_error: BiocParallel errors
##  1 remote errors, element index: 2
##  3 unevaluated and other errors
##  first remote error:
## Error in FUN(...): non-numeric argument to mathematical function
## >
## results and errors available as 'bpresult(x)'
```

Errors, Logs and Debugging in *BiocParallel*

```
bpresult(result)

## [[1]]
## [1] 1
##
## [[2]]
## <remote_error in FUN(...): non-numeric argument to mathematical function>
## traceback() available as 'attr(x, "traceback")'
##
## [[3]]
## [1] 1.732051
##
## [[4]]
## <unevaluated_error: not evaluated due to previous error>
##
## [[5]]
## <unevaluated_error: not evaluated due to previous error>
##
## [[6]]
## <unevaluated_error: not evaluated due to previous error>
##
## attr(,"REDOENV")
## <environment: 0x560a5c62e910>
```

Using `stop.on.error=FALSE`, all tasks are evaluated.

```
X <- list("1", 2, 3, 4, 5, 6)
param <- SnowParam(3, tasks = length(X), stop.on.error = FALSE)
result <- tryCatch({
  bplapply(X, sqrt, BPPARAM = param)
}, error=identity)
result

## <bplist_error: BiocParallel errors
## 1 remote errors, element index: 1
## 0 unevaluated and other errors
## first remote error:
## Error in FUN(...): non-numeric argument to mathematical function
## >
## results and errors available as 'bpresult(x)'
```

```
bpresult(result)

## [[1]]
## <remote_error in FUN(...): non-numeric argument to mathematical function>
## traceback() available as 'attr(x, "traceback")'
##
## [[2]]
## [1] 1.414214
##
## [[3]]
## [1] 1.732051
##
```

Errors, Logs and Debugging in *BiocParallel*

```
## [[4]]
## [1] 2
##
## [[5]]
## [1] 2.236068
##
## [[6]]
## [1] 2.44949
##
## attr(,"REDOENV")
## <environment: 0x560a5c4f21c0>
```

`bpttry()` is a convenient way of trying to evaluate a `bpapply`-like expression, returning the evaluated results without signalling an error.

```
bpttry({
  bpapply(X, sqrt, BPPARAM=param)
})
## [[1]]
## <remote_error in FUN(...): non-numeric argument to mathematical function>
## traceback() available as 'attr(x, "traceback")'
##
## [[2]]
## [1] 1.414214
##
## [[3]]
## [1] 1.732051
##
## [[4]]
## [1] 2
##
## [[5]]
## [1] 2.236068
##
## [[6]]
## [1] 2.44949
##
## attr(,"REDOENV")
## <environment: 0x560a5c5053c0>
```

In the next example the elements of `X` are grouped instead of run separately. The default value for `tasks` is 0 which means 'X' is split as evenly as possible across the number of workers. There are 3 workers so the first task consists of `list(1, 2)`, the second is `list("3", 4)` and the third is `list(5, 6)`.

```
X <- list(1, 2, "3", 4, 5, 6)
param <- SnowParam(3, stop.on.error = TRUE)
```

The output shows an error in when evaluating the third element, but also that the fourth element, in the same chunk as 3, was not evaluated. All elements are evaluated because they were assigned to workers before the first error occurred.

```
bptry(bplapply(X, sqrt, BPPARAM = param))

## [[1]]
## [1] 1
##
## [[2]]
## [1] 1.414214
##
## [[3]]
## <remote_error in FUN(...): non-numeric argument to mathematical function>
## traceback() available as 'attr(x, "traceback")'
##
## [[4]]
## <unevaluated_error: not evaluated due to previous error>
##
## [[5]]
## [1] 2.236068
##
## [[6]]
## [1] 2.44949
##
## attr(,"REDOENV")
## <environment: 0x560a5c621e20>
```

Side Note: Results are collected from workers as they finish which is not necessarily the same order in which they were loaded. Depending on how tasks are divided it is possible that the task with the error completes after all others so essentially all workers complete before the job is stopped. In this situation the output includes all results along with the error message and it may appear that `stop.on.error=TRUE` did not stop the job soon enough. This is just a heads up that the usefulness of `stop.on.error=TRUE` may vary with run time and distribution of tasks over workers.

2.3 Identify failures with `bpok()`

The `bpok()` function is a quick way to determine which (if any) tasks failed. In this example we use `bptry()` to retrieve the partially evaluated expression, including the failed elements.

```
param <- SnowParam(2, stop.on.error=FALSE)
result <- bptry(bplapply(list(1, "2", 3), sqrt, BPPARAM=param))
```

`bpok` returns TRUE if the task was successful.

```
bpok(result)
## [1] TRUE FALSE TRUE
```

Once errors are identified with `bpok` the traceback can be retrieved with the `attr` function. This is possible because errors are returned as `condition` objects with the traceback as an attribute.

```
attr(result[[which(!bpok(result))]], "traceback")
## [1] "3: handle_error(e)"
```

Errors, Logs and Debugging in *BiocParallel*

```
## [2] "2: h(simpleError(msg, call))"
## [3] "1: .handleSimpleError(function (e) "
## [4] "  {"
## [5] "      annotated_condition <- handle_error(e)"
## [6] "      stop(annotated_condition)"
## [7] "  }, \"non-numeric argument to mathematical function\", base::quote(FUN(...)))"
```

Note that the traceback has been modified from the full traceback provided by *R* to include only the calls from the time the `bplapply` FUN is evaluated.

2.4 Rerun failed tasks with `BPRED0`

Tasks can fail due to hardware problems or bugs in the input data. The *BiocParallel* functions support a `BPRED0` (re-do) argument for recomputing only the tasks that failed. A list of partial results and errors is supplied to `BPRED0` in a second call to the function. The failed elements are identified, recomputed and inserted into the original results.

The bug in this example is the second element of 'X' which is a character when it should be numeric.

```
X <- list(1, "2", 3)
param <- SnowParam(2, stop.on.error=FALSE)
result <- bptry(bplapply(X, sqrt, BPPARAM=param))
result

## [[1]]
## [1] 1
##
## [[2]]
## <remote_error in FUN(...): non-numeric argument to mathematical function>
## traceback() available as 'attr(x, "traceback")'
##
## [[3]]
## [1] 1.732051
##
## attr(,"REDOENV")
## <environment: 0x560a5c693ae8>
```

First fix the input data.

```
X.redo <- list(1, 2, 3)
```

Repeat the call to `bplapply` this time supplying the partial results as `BPRED0`. Only the failed calculations are computed, in the present case requiring only one worker.

```
bplapply(X.redo, sqrt, BPRED0=result, BPPARAM=param)

## [[1]]
## [1] 1
##
## [[2]]
## [1] 1.414214
##
```

```
## [[3]]  
## [1] 1.732051
```

3 Logging

NOTE: Logging as described in this section is supported for `SnowParam`, `MulticoreParam` and `SerialParam`.

3.1 Parameters

Logging in *BiocParallel* is controlled by 3 fields in the `BiocParallelParam`:

```
log:      TRUE or FALSE  
logdir:   location to write log file  
threshold: one of "TRACE", "DEBUG", "INFO", "WARN", "ERROR", "FATAL"
```

When `log = TRUE` the *futile.logger* package is loaded on each worker. *BiocParallel* uses a custom script on the workers to collect log messages as well as additional statistics such as `gc`, runtime and node information. Output to `stderr` and `stdout` is also captured.

By default `log` is `FALSE` and `threshold` is `INFO`.

```
param <- SnowParam(stop.on.error=FALSE)  
param  
  
## class: SnowParam  
##  bpisup: FALSE; bpnworkers: 2; bptasks: 0; bpjobname: BPJOB  
##  bplog: FALSE; bpthreshold: INFO; bpstopOnError: FALSE  
##  bpRNGseed: ; bptimeout: NA; bpprogressbar: FALSE  
##  bpexportglobals: TRUE; bpexportvariables: TRUE; bpforceGC: FALSE  
##  bpfallback: TRUE  
##  bplogdir: NA  
##  bresultdir: NA  
##  cluster type: SOCK
```

Turn logging on and set the threshold to `TRACE`.

```
bplog(param) <- TRUE  
bpthreshold(param) <- "TRACE"  
param  
  
## class: SnowParam  
##  bpisup: FALSE; bpnworkers: 2; bptasks: 0; bpjobname: BPJOB  
##  bplog: TRUE; bpthreshold: TRACE; bpstopOnError: FALSE  
##  bpRNGseed: ; bptimeout: NA; bpprogressbar: FALSE  
##  bpexportglobals: TRUE; bpexportvariables: TRUE; bpforceGC: FALSE  
##  bpfallback: TRUE  
##  bplogdir: NA  
##  bresultdir: NA  
##  cluster type: SOCK
```


3.2 Setting a threshold

All thresholds defined in *futile.logger* are supported: *FATAL*, *ERROR*, *WARN*, *INFO*, *DEBUG* and *TRACE*. All messages greater than or equal to the severity of the threshold are shown. For example, a threshold of *INFO* will print all messages tagged as *FATAL*, *ERROR*, *WARN* and *INFO*.

Because the default threshold is *INFO* it catches the *ERROR*-level message thrown when attempting the square root of a character ("2").

```
tryCatch({
  bplapply(list(1, "2", 3), sqrt, BPPARAM = param)
}, error=function(e) invisible(e))

## ##### LOG OUTPUT #####
## Task: 2
## Node: 1
## Timestamp: 2024-06-25 05:05:07.674168
## Success: TRUE
##
## Task duration:
##   user system elapsed
## 0.041 0.002 0.042
##
## Memory used:
##      used (Mb) gc trigger (Mb) max used (Mb)
## Ncells 7308399 390.4 12246307 654.1 8601920 459.4
## Vcells 13010186 99.3 21929244 167.4 15105743 115.3
##
## Log messages:
## INFO [2024-06-25 05:05:07] loading futile.logger package
##
## stderr and stdout:

## ##### LOG OUTPUT #####
## Task: 1
## Node: 2
## Timestamp: 2024-06-25 05:05:08.136384
## Success: FALSE
##
## Task duration:
##   user system elapsed
## 0.044 0.002 0.046
##
## Memory used:
##      used (Mb) gc trigger (Mb) max used (Mb)
## Ncells 7309013 390.4 12246307 654.1 8601920 459.4
## Vcells 13011613 99.3 21929244 167.4 15105743 115.3
##
## Log messages:
## INFO [2024-06-25 05:05:07] loading futile.logger package
```

Errors, Logs and Debugging in *BiocParallel*

```
## ERROR [2024-06-25 05:05:07] non-numeric argument to mathematical function
##
## stderr and stdout:
```

All user-supplied messages written in the *futile.logger* syntax are also captured. This function performs argument checking and includes a couple of *WARN* and *DEBUG*-level messages.

```
FUN <- function(i) {
  futile.logger::flog.debug(paste("value of 'i':", i))

  if (!length(i)) {
    futile.logger::flog.warn("'i' has length 0")
    NA
  } else if (!is(i, "numeric")) {
    futile.logger::flog.debug("coercing 'i' to numeric")
    as.numeric(i)
  } else {
    i
  }
}
```

Turn logging on and set the threshold to *WARN*.

```
param <- SnowParam(2, log = TRUE, threshold = "WARN", stop.on.error=FALSE)
result <- bplapply(list(1, "2", integer()), FUN, BPPARAM = param)
```

```
## ##### LOG OUTPUT #####
## Task: 1
## Node: 2
## Timestamp: 2024-06-25 05:05:11.294208
## Success: TRUE
##
## Task duration:
##   user system elapsed
## 1.052 0.041 1.093
##
## Memory used:
##           used (Mb) gc trigger (Mb) max used (Mb)
## Ncells 7310096 390.5 12246307 654.1 8601920 459.4
## Vcells 13014564 99.3 21929244 167.4 15105743 115.3
##
## Log messages:
##
##
## stderr and stdout:
## ##### LOG OUTPUT #####
## Task: 2
## Node: 1
## Timestamp: 2024-06-25 05:05:11.760481
## Success: TRUE
##
## Task duration:
```

Errors, Logs and Debugging in *BiocParallel*

```
## user system elapsed
## 1.081 0.027 1.108
##
## Memory used:
## used (Mb) gc trigger (Mb) max used (Mb)
## Ncells 7310119 390.5 12246307 654.1 8601920 459.4
## Vcells 13014623 99.3 21929244 167.4 15105743 115.3
##
## Log messages:
## WARN [2024-06-25 05:05:11] 'i' has length 0
##
## stderr and stdout:
simplify2array(result)
```

Changing the threshold to *DEBUG* catches both *WARN* and *DEBUG* messages.

```
param <- SnowParam(2, log = TRUE, threshold = "DEBUG", stop.on.error=FALSE)
result <- bplapply(list(1, "2", integer()), FUN, BPPARAM = param)

## ##### LOG OUTPUT #####
## Task: 1
## Node: 2
## Timestamp: 2024-06-25 05:05:14.939363
## Success: TRUE
##
## Task duration:
## user system elapsed
## 1.057 0.043 1.100
##
## Memory used:
## used (Mb) gc trigger (Mb) max used (Mb)
## Ncells 7310034 390.4 12246307 654.1 8601920 459.4
## Vcells 13014800 99.3 21929244 167.4 15105743 115.3
##
## Log messages:
## INFO [2024-06-25 05:05:13] loading futile.logger package
## DEBUG [2024-06-25 05:05:14] value of 'i': 1
## DEBUG [2024-06-25 05:05:14] value of 'i': 2
## DEBUG [2024-06-25 05:05:14] coercing 'i' to numeric
##
## stderr and stdout:

## ##### LOG OUTPUT #####
## Task: 2
## Node: 1
## Timestamp: 2024-06-25 05:05:15.409657
## Success: TRUE
##
## Task duration:
## user system elapsed
```

```
## 1.049 0.059 1.109
##
## Memory used:
##      used (Mb) gc trigger (Mb) max used (Mb)
## Ncells 7310055 390.4 12246307 654.1 8601920 459.4
## Vcells 13014867 99.3 21929244 167.4 15105743 115.3
##
## Log messages:
## INFO [2024-06-25 05:05:13] loading futile.logger package
## DEBUG [2024-06-25 05:05:14] value of 'i':
## WARN [2024-06-25 05:05:14] 'i' has length 0
##
## stderr and stdout:
simplify2array(result)
```

3.3 Log files

When `log == TRUE`, log messages are written to the console by default. If `logdir` is given the output is written out to files, one per task. File names are prefixed with the name in `bpjobname(BPPARAM)`; default is 'BPJOB'.

```
param <- SnowParam(2, log = TRUE, threshold = "DEBUG", logdir = tempdir())
res <- bplapply(list(1, "2", integer()), FUN, BPPARAM = param)
## loading futile.logger on workers
list.files(bplogdir(param))
## [1] "BPJOB.task1.log" "BPJOB.task2.log"
```

Read in BPJOB.task2.log:

```
readLines(paste0(bplogdir(param), "/BPJOB.task2.log"))

## [1] "##### LOG OUTPUT #####"
## [2] "Task: 2"
## [3] "Node: 2"
## [4] "Timestamp: 2015-07-08 09:03:59"
## [5] "Success: TRUE"
## [6] "Task duration: "
## [7] " user system elapsed "
## [8] " 0.009 0.000 0.011 "
## [9] "Memory use (gc): "
## [10] "      used (Mb) gc trigger (Mb) max used (Mb)"
## [11] "Ncells 325664 17.4 592000 31.7 393522 21.1"
## [12] "Vcells 436181 3.4 1023718 7.9 530425 4.1"
## [13] "Log messages:"
## [14] "DEBUG [2015-07-08 09:03:59] value of 'i': 2"
## [15] "INFO [2015-07-08 09:03:59] coercing to numeric"
## [16] "DEBUG [2015-07-08 09:03:59] value of 'i': "
## [17] "WARN [2015-07-08 09:03:59] 'i' is missing"
## [18] ""
## [19] "stderr and stdout:"
```

```
## [20] "character(0)"
```

4 Worker timeout

NOTE: `timeout` is supported for `SnowParam` and `MulticoreParam`.

For long running jobs or untested code it can be useful to set a time limit. The `timeout` field is the time, in seconds, allowed for each worker to complete a task; default is `Inf`. If the task takes longer than `timeout` a timeout error is returned.

Time can be changed during param construction with the `timeout` arg,

```
param <- SnowParam(timeout = 20, stop.on.error=FALSE)
param

## class: SnowParam
##  bpisup: FALSE; bpnworkers: 2; bptasks: 0; bpjobname: BPJOB
##  bplog: FALSE; bpthreshold: INFO; bpstopOnError: FALSE
##  bpRNGseed: ; bptimeout: 20; bpprogressbar: FALSE
##  bpexportglobals: TRUE; bpexportvariables: TRUE; bpforceGC: FALSE
##  bpfallback: TRUE
##  bplogdir: NA
##  bpresultdir: NA
##  cluster type: SOCK
```

or with the `bptimeout` setter:

```
param <- SnowParam(timeout = 2, stop.on.error=FALSE)
fun <- function(i) {
  Sys.sleep(i)
  i
}
bpttry(bplapply(1:3, fun, BPPARAM = param))

## [[1]]
## [1] 1
##
## [[2]]
## [1] 2
##
## [[3]]
## [1] 3
```

5 Debugging

Effective debugging strategies vary by problem and often involve a combination of error handling and logging techniques. In general, when debugging *R*-generated errors the traceback is often the best place to start followed by adding debug messages to the worker function. When trouble shooting unexpected behavior (i.e., not a formal error or warning) adding debug messages or switching to `SerialParam` are good approaches. Below is an overview of these different strategies.

5.1 Accessing the traceback

The traceback is a good place to start when tracking down *R*-generated errors. Because the function is executed on the workers it's not accessible for interactive debugging with functions such as `trace` or `debug`. The traceback provides a snapshot of the state of the worker at the time the error was thrown.

This function takes the square root of the absolute value of a vector.

```
fun1 <- function(x) {
  v <- abs(x)
  sapply(1:length(v), function(i) sqrt(v[i]))
}
```

Calling "fun1" with a character throws an error:

```
param <- SnowParam(stop.on.error=FALSE)
result <- bpttry({
  bplapply(list(c(1,3), 5, "6"), fun1, BPPARAM = param)
})
result
## [[1]]
## [1] 1.000000 1.732051
##
## [[2]]
## [1] 2.236068
##
## [[3]]
## <remote_error in abs(x): non-numeric argument to mathematical function>
## traceback() available as 'attr(x, "traceback")'
##
## attr(,"REDOENV")
## <environment: 0x11bdb3a18>
```

Identify which elements failed with `bpok`:

```
bpok(result)

## [1] TRUE TRUE FALSE
```

The error (i.e., third element of "res") is a `condition` object:

```
is(result[[3]], "condition")

## [1] TRUE
```

The traceback is an attribute of the `condition` and can be accessed with the `attr` function.

```
cat(attr(result[[3]], "traceback"), sep = "\n")
## 4: handle_error(e)
## 3: h(simpleError(msg, call))
## 2: .handleSimpleError(function (e)
##   {
##     annotated_condition <- handle_error(e)
##     stop(annotated_condition)
##   }, "non-numeric argument to mathematical function", base::quote(abs(x))) at #2
```

```
## 1: FUN(...)
```

In this example the error occurs in `FUN`; lines 2, 3, 4 involve error handling.

5.2 Adding debug messages

When a `numeric()` is passed to “fun1” no formal error is thrown but the length of the second list element is 2 when it should be 1.

```
bplapply(list(c(1,3), numeric(), 6), fun1, BPPARAM = param)
```

```
## [[1]]
## [1] 1.000000 1.732051
##
## [[2]]
## [[2]][[1]]
## [1] NA
##
## [[2]][[2]]
## numeric(0)
##
## [[3]]
## [1] 2.44949
```

Without a formal error we have no traceback so we'll add a few debug messages. The `futile.logger` syntax tags messages with different levels of severity. A message created with `flog.debug` will only print if the threshold is `DEBUG` or lower. So in this case it will catch both `INFO` and `DEBUG` messages.

“fun2” has debug statements that show the value of ‘x’, length of ‘v’ and the index ‘i’.

```
fun2 <- function(x) {
  v <- abs(x)
  futile.logger::flog.debug(
    paste0("'x' = ", paste(x, collapse=","), ": length(v) = ", length(v))
  )
  sapply(1:length(v), function(i) {
    futile.logger::flog.info(paste0("'i' = ", i))
    sqrt(v[i])
  })
}
```

Create a param that logs at a threshold level of `DEBUG`.

```
param <- SnowParam(3, log = TRUE, threshold = "DEBUG")
```

```
res <- bplapply(list(c(1,3), numeric(), 6), fun2, BPPARAM = param)

## ##### LOG OUTPUT #####
## Task: 1
## Node: 3
## Timestamp: 2024-06-25 05:05:25.051815
## Success: TRUE
##
```

Errors, Logs and Debugging in *BiocParallel*

```
## Task duration:
##   user system elapsed
##  1.132  0.019  1.150
##
## Memory used:
##           used (Mb) gc trigger (Mb) max used (Mb)
## Ncells  7310344 390.5  12246307 654.1  8601920 459.4
## Vcells 13017156  99.4   21929244 167.4 15105743 115.3
##
## Log messages:
## INFO [2024-06-25 05:05:23] loading futile.logger package
## DEBUG [2024-06-25 05:05:25] 'x' = 1,3: length(v) = 2
## INFO [2024-06-25 05:05:25] 'i' = 1
## INFO [2024-06-25 05:05:25] 'i' = 2
##
## stderr and stdout:
## ##### LOG OUTPUT #####
## Task: 2
## Node: 2
## Timestamp: 2024-06-25 05:05:25.537727
## Success: TRUE
##
## Task duration:
##   user system elapsed
##  1.517  0.055  1.577
##
## Memory used:
##           used (Mb) gc trigger (Mb) max used (Mb)
## Ncells  7310372 390.5  12246307 654.1  8601920 459.4
## Vcells 13017234  99.4   21929244 167.4 15105743 115.3
##
## Log messages:
## INFO [2024-06-25 05:05:23] loading futile.logger package
## DEBUG [2024-06-25 05:05:25] 'x' = : length(v) = 0
## INFO [2024-06-25 05:05:25] 'i' = 1
## INFO [2024-06-25 05:05:25] 'i' = 0
##
## stderr and stdout:
## ##### LOG OUTPUT #####
## Task: 3
## Node: 1
## Timestamp: 2024-06-25 05:05:26.001379
## Success: TRUE
##
## Task duration:
##   user system elapsed
##  1.615  0.034  1.652
##
## Memory used:
##           used (Mb) gc trigger (Mb) max used (Mb)
```


Errors, Logs and Debugging in *BiocParallel*

```
## Ncells 7310393 390.5 12246307 654.1 8601920 459.4
## Vcells 13017303 99.4 21929244 167.4 15105743 115.3
##
## Log messages:
## INFO [2024-06-25 05:05:23] loading futile.logger package
## DEBUG [2024-06-25 05:05:25] 'x' = 6: length(v) = 1
## INFO [2024-06-25 05:05:25] 'i' = 1
##
## stderr and stdout:

res
```

The debug messages require close inspection, but focusing on task 2 we see

```
res
## ##### LOG OUTPUT #####
## Task: 2
## Node: 2
## Timestamp: 2023-03-23 12:17:28.969158
## Success: TRUE
##
## Task duration:
##   user  system elapsed
## 0.156  0.005  0.163
##
## Memory used:
##      used (Mb) gc trigger (Mb) limit (Mb) max used (Mb)
## Ncells 942951 50.4 1848364 98.8 NA 1848364 98.8
## Vcells 1941375 14.9 8388608 64.0 32768 2446979 18.7
##
## Log messages:
## INFO [2023-03-23 12:17:28] loading futile.logger package
## DEBUG [2023-03-23 12:17:28] 'x' = : length(v) = 0
## INFO [2023-03-23 12:17:28] 'i' = 1
## INFO [2023-03-23 12:17:28] 'i' = 0
##
## stderr and stdout:
```

This reveals the problem. The index for `sapply` is along 'v' which in this case has length 0. This forces 'i' to take values of '1' and '0' giving an output of length 2 for the second element (i.e., NA and `numeric(0)`).

"fun2" can be fixed by using `seq_along(v)` to create the index instead of `1:length(v)`.

5.3 Local debugging with `SerialParam`

Errors that occur on parallel workers can be difficult to debug. Often the traceback sent back from the workers is too much to parse or not informative. We are also limited in that our interactive strategies of `browser` and `trace` are not available.

One option for further debugging is to run the code in serial with `SerialParam`. This removes the "parallel" component and is the same as running a straight `*apply` function. This approach may not help if the problem was hardware related but can be very useful when the bug is in the R code.

We use the now familiar square root example with a bug in the second element of X.

```
res <- bptry({
  bplapply(list(1, "2", 3), sqrt,
           BPPARAM = SnowParam(3, stop.on.error=FALSE))
})
result
## [[1]]
## [1] 1
##
## [[2]]
## [1] 2
##
## [[3]]
## [1] NA
```

`sqrt` is an internal function. The problem is likely with our data going into the function and not the `sqrt` function itself. We can write a small wrapper around `sqrt` so we can see the input.

```
fun3 <- function(i) sqrt(i)
```

Debug the new function:

```
debug(fun3)
```

We want to recompute only elements that failed and for that we use the `BPREDO` argument. The `BPPARAM` has been changed to `SerialParam` so the job is run in the local workspace in serial.

```
> bplapply(list(1, "2", 3), fun3, BPREDO = result, BPPARAM = SerialParam())
Resuming previous calculation ...
debugging in: FUN(...)
debug: sqrt(i)
Browse[2]> objects()
[1] "i"
Browse[2]> i
[1] "2"
```

Browse[2]>

The local browsing allowed us to see the problem input was the character "2".

6 sessionInfo()

`toLatex(sessionInfo())`

- R version 4.4.1 (2024-06-14), x86_64-pc-linux-gnu
- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_US.UTF-8, LC_COLLATE=C, LC_MONETARY=en_US.UTF-8, LC_MESSAGES=en_US.UTF-8, LC_PAPER=en_US.UTF-8, LC_NAME=C, LC_ADDRESS=C, LC_TELEPHONE=C, LC_MEASUREMENT=en_US.UTF-8, LC_IDENTIFICATION=C
- Time zone: Etc/UTC
- TZcode source: system (glibc)
- Running under: Ubuntu 24.04 LTS
- Matrix products: default
- BLAS: /usr/lib/x86_64-linux-gnu/openblas-pthread/libblas.so.3
- LAPACK: /usr/lib/x86_64-linux-gnu/openblas-pthread/libopenblas-p-r0.3.26.so ; LAPACK version3.12.0
- Base packages: base, datasets, grDevices, graphics, methods, stats, stats4, utils
- Other packages: Biobase 2.65.0, BiocGenerics 0.51.0, BiocParallel 1.39.0, BiocStyle 2.33.1, Biostrings 2.73.1, GenomelInfoDb 1.41.1, GenomicAlignments 1.41.0, GenomicRanges 1.57.1, IRanges 2.39.0, MatrixGenerics 1.17.0, RNAseqData.HNRNPC.bam.chr14 0.43.0, Rsamtools 2.21.0, S4Vectors 0.43.0, SummarizedExperiment 1.35.0, XVector 0.45.0, matrixStats 1.3.0
- Loaded via a namespace (and not attached): BiocManager 1.30.23, DelayedArray 0.31.3, GenomelInfoDbData 1.2.12, Matrix 1.7-0, R6 2.5.1, S4Arrays 1.5.1, SparseArray 1.5.10, UCSC.utils 1.1.0, abind 1.4-5, backports 1.5.0, base64url 1.4, batchtools 0.9.17, bitops 1.0-7, brew 1.0-10, bslib 0.7.0, buildtools 1.0.0, cachem 1.1.0, checkmate 2.3.1, cli 3.6.3, codetools 0.2-20, compiler 4.4.1, crayon 1.5.3, data.table 1.15.4, digest 0.6.36, evaluate 0.24.0, fastmap 1.2.0, fs 1.6.4, grid 4.4.1, highr 0.11, hms 1.1.3, htmltools 0.5.8.1, httr 1.4.7, jquerylib 0.1.4, jsonlite 1.8.8, knitr 1.47, lattice 0.22-6, lifecycle 1.0.4, maketools 1.3.0, parallel 4.4.1, pkgconfig 2.0.3, prettyunits 1.2.0, progress 1.2.3, rappdirs 0.3.3, rlang 1.1.4, rmarkdown 2.27, sass 0.4.9, snow 0.4-4, stringi 1.8.4, sys 3.4.2, tools 4.4.1, vctrs 0.6.5, withr 3.0.0, xfun 0.45, yaml 2.3.8, zlibbioc 1.51.1