Errors, Logs and Debugging in *BiocParallel*

Valerie Obenchain and Martin Morgan

Edited: May 13, 2015; Compiled: June 23, 2015

Contents

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*, *BatchJobs* 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 handling existed, such as *BatchJobs* and *foreach*, those behaviors were preserved. Clusters created with *snow* and *parallel* now have flexible error handling and logging available through `SnowParma` 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 `bplapply`). 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 `biocLite`:

```
source("http://bioconductor.org/biocLite.R")
biocLite("BiocParallel")
```

Load the package:

```
library(BiocParallel)
```

2 Error Handling

2.1 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 `FALSE` by default.

```
param <- SnowParam()
param
```

```
## class: SnowParam
##   bpworkers:6; bptasks:0; bpRNGseed:; bpisup:FALSE
##   bplog:FALSE; bpthreshold:INFO; bplogdir:NA
##   bpstopOnError:FALSE; bpprogressbar:FALSE
##   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
##   bpworkers:2; bptasks:0; bpRNGseed:; bpisup:FALSE
##   bplog:FALSE; bpthreshold:INFO; bplogdir:NA
##   bpstopOnError:TRUE; bpprogressbar:FALSE
##   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. To more clearly demonstrate the affect of `stop.on.error` the number of tasks is set equal to the length of `X`. This forces each element of `X` to be executed separately (6 tasks) vs chunked.

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

The output list contains results for tasks 1 and 2 and an error for task 3. Tasks 4, 5, and 6 are not attempted.

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

```
## Warning in bpdynamicClusterApply(bpbackend(BPPARAM), lapply, length(X), :  error in task 3
```

```
## [[1]]
## [1] 1
##
## [[2]]
## [1] 1.414214
##
## [[3]]
## <remote-error in FUN(...): non-numeric argument to mathematical function>
```

Next we look at an example where 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)
```

To simulate a longer running computation sleep time is added to 'FUN'. The sleep forces task 2 to finish before task 3.

```
FUN <- function(i) { Sys.sleep(i); sqrt(i) }
```

The output shows an error in task 2 (vs 3 in the previous example) and a result for '4' is included because it was part of the second task.

```
bplapply(X, FUN, BPPARAM = param)
```

```
## Warning in bpdynamicClusterApply(bpbackend(BPPARAM), lapply, length(X), :  error in task 2
```

```
## [[1]]
## [1] 1
##
## [[2]]
## [1] 1.414214
##
## [[3]]
## <remote-error in sqrt(i): non-numeric argument to mathematical function>
##
## [[4]]
## [1] 2
```

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 among workers it is possible that the task with an error completes after all others. In that situation the output will includes all results along with the error message and it may appear that `stop.on.error` is not doing much good. This is simply a heads up that the usefulness of `stop.on.error` may vary with run time and distribution of tasks over workers.

2.2 Identify failures with `bpok`

The `bpok()` function is a quick way to determine which (if any) tasks failed. In this example the second element fails.

```
res <- bplapply(list(1, "2", 3), sqrt)
res
```

```
## [[1]]
## [1] 1
##
## [[2]]
## <remote-error in FUN(...): non-numeric argument to mathematical function>
##
## [[3]]
## [1] 1.732051
```

bpok returns TRUE if the task was successful.

```
bpok(res)
```

```
## [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.

```
fail <- !bpok(res)
tail(attr(res[[2]], "traceback"))
```

```
## [1] "       call <- sapply(sys.calls(), deparse)"
## [2] "       e <- structure(e, class = c(\"remote-error\", \"condition\"), "
## [3] "           traceback = capture.output(traceback(call)))"
## [4] "       invokeRestart(\"abort\", e)"
## [5] "   }, \"non-numeric argument to mathematical function\", quote(FUN(...)))"
## [6] "1: h(simpleError(msg, call))"
```

2.3 Rerun failed tasks with `BPREDO`

Tasks can fail due to hardware problems or bugs in the input data. The *BiocParallel* functions support a `BPREDO` (re-do) argument for recomputing only the tasks that failed. A list of partial results and errors is supplied to `BPREDO` 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)
res <- bplapply(X, sqrt)
res
```

```
## [[1]]
## [1] 1
##
## [[2]]
## <remote-error in FUN(...): non-numeric argument to mathematical function>
##
## [[3]]
## [1] 1.732051
```

First fix the input data.

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

Repeat the call to `bplapply` this time supplying the partial results as `BPREDO`.

```
bplapply(X.redo, sqrt, BPREDO = res)
```

```
## Resuming previous calculation ...
```

```
## [[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.

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

```
param <- SnowParam()
param
```

```
## class: SnowParam
##   bpworkers:6; bptasks:0; bpRNGseed:; bpisup:FALSE
##   bplog:FALSE; bpthreshold:INFO; bplogdir:NA
##   bpstopOnError:FALSE; bpprogressbar:FALSE
##   bpresultdir:NA
## cluster type: SOCK
```

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

```
bplog(param) <- TRUE
bpthreshold(param) <- "TRACE"
param
```

```
## class: SnowParam
##   bpworkers:6; bptasks:0; bpRNGseed:; bpisup:FALSE
##   bplog:TRUE; bpthreshold:TRACE; bplogdir:NA
##   bpstopOnError:FALSE; bpprogressbar:FALSE
##   bpresultdir:NA
## cluster type: SOCK
```

3.2 Setting a threshold

All thresholds defined in *futile.logger* are all 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*.

In this code chunk an *INFO*-level message is emitted when *futile.logger* is loaded on the workers and a *ERROR*-level message when attempting the square root of a character ("2").

```
bpthreshold(param) <- "INFO"
bplapply(list(1, "2", 3), sqrt, BPPARAM = param)
```

```
## INFO [2015-06-23 20:47:33] loading futile.logger on workers
## [1] "ERROR [2015-06-23 20:47:33] non-numeric argument to mathematical function\n"
## [[1]]
## [1] 1
##
## [[2]]
## <remote-error in FUN(...): non-numeric argument to mathematical function>
##
## [[3]]
## [1] 1.732051
```

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) {
  flog.debug(paste0("value of 'i': ", i))

  if (!length(i)) {
      flog.warn("'i' is missing")
      NA
  } else if (!is(i, "numeric")) {
      flog.warn("coercing to numeric")
      as.numeric(i)
  } else {
      i
  }
}
```

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

```
param <- SnowParam(2, log = TRUE, threshold = "WARN")
bplapply(list(1, "2", integer()), FUN, BPPARAM = param)
```

```
## [1] "WARN [2015-06-23 20:47:34] coercing to numeric\n"
## [2] "WARN [2015-06-23 20:47:34] 'i' is missing\n"
## [[1]]
## [1] 1
##
## [[2]]
## [1] 2
##
## [[3]]
## [1] NA
```

Changing the threshold to *DEBUG* catches all *WARN*, *INFO* and *DEBUG* messages.

```
param <- SnowParam(2, log = TRUE, threshold = "DEBUG")
bplapply(list(1, "2", integer()), FUN, BPPARAM = param)
```

```
## INFO [2015-06-23 20:47:35] loading futile.logger on workers
## [1] "DEBUG [2015-06-23 20:47:35] value of 'i': 1\n"
## [1] "DEBUG [2015-06-23 20:47:35] value of 'i': 2\n"
## [2] "WARN [2015-06-23 20:47:35] coercing to numeric\n"
## [3] "DEBUG [2015-06-23 20:47:35] value of 'i': \n"
## [4] "WARN [2015-06-23 20:47:35] 'i' is missing\n"
## [[1]]
## [1] 1
##
## [[2]]
## [1] 2
##
## [[3]]
## [1] NA
```

3.3 Log files

futile.logger messages are written to the console by default. When `logdir` is given, these messages along with additional

statistics of memory use, duration, node name etc. are written to a file called "BPLOG.out".

```
> param <- SnowParam(2, log = TRUE, threshold = "DEBUG", logdir = tempdir())
> res <- bplapply(list(1, "2", integer()), FUN, BPPARAM = param)
> list.files(bplogdir(param))
[1] "BPLOG.out"
```

Read in the log file from `logdir`:

```
> readLines(paste0(bplogdir(param), "/BPLOG.out"))
 [1] ################ LOG OUTPUT ################
 [2] Task: 2
 [3] Node: 2
 [4] Timestamp: 2015-05-15 07:29:11
 [5] Success: TRUE
 [6] Task duration:
 [7]    user  system elapsed
 [8]   0.011   0.000   0.011
 [9] Memory use (gc):
[10]          used (Mb) gc trigger (Mb) max used (Mb)
[11] Ncells 325582 17.4     592000 31.7   382191 20.5
[12] Vcells 435234  3.4    1023718  7.9   526789  4.1
[13] Log messages:
[14] DEBUG [2015-05-15 07:29:11] value of 'i': 2
[15] WARN [2015-05-15 07:29:11] coercing to numeric
[16] DEBUG [2015-05-15 07:29:11] value of 'i':
[17] WARN [2015-05-15 07:29:11] 'i' is missing
[18]
[19] ################ LOG OUTPUT ################
[20] Task: 1
[21] Node: 1
[22] Timestamp: 2015-05-15 07:29:11
[23] Success: TRUE
[24] Task duration:
[25]    user  system elapsed
[26]   0.006   0.000   0.005
[27] Memory use (gc):
[28]          used (Mb) gc trigger (Mb) max used (Mb)
[29] Ncells 325547 17.4     592000 31.7   341650 18.3
[30] Vcells 435185  3.4    1023718  7.9   517996  4.0
[31] Log messages:
[32] DEBUG [2015-05-15 07:29:11] value of 'i': 1
```

4 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.

4.1 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:

```
res <- bplapply(list(c(1,3), 5, "6"), fun1)
res
```

```
## [[1]]
## [1] 1.000000 1.732051
##
## [[2]]
## [1] 2.236068
##
## [[3]]
## <remote-error in abs(x): non-numeric argument to mathematical function>
```

Identify which elements failed with `bpok`:

```
bpok(res)
```

```
## [1]  TRUE  TRUE FALSE
```

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

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

```
## [1] TRUE
```

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

```
noquote(tail(attr(res[[3]], "traceback")))
```

```
## [1]         call <- sapply(sys.calls(), deparse)
## [2]         e <- structure(e, class = c("remote-error", "condition"),
## [3]             traceback = capture.output(traceback(call)))
## [4]         invokeRestart("abort", e)
## [5]     }, "non-numeric argument to mathematical function", quote(abs(x)))
## [6] 1: h(simpleError(msg, call))
```

4.2 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)
```

```
## [[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 try adding a few debug messages to "fun1". 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.

"fun2" has debug statements that show the value of 'x', length of 'v' and the index 'i'.

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

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

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

The debug messages reveal the problem occurs when 'x' is `numeric()`. 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)`).

```
param <- SnowParam(3, log = TRUE, threshold = "DEBUG")
res <- bplapply(list(c(1,3), numeric(), 6), fun2, BPPARAM = param)
```

```
## INFO [2015-06-23 20:47:36] loading futile.logger on workers
## [1] "DEBUG [2015-06-23 20:47:36] 'x' = 6: length(v) = 1\n"
## [2] "DEBUG [2015-06-23 20:47:36] 'i': 1\n"
## [1] "DEBUG [2015-06-23 20:47:36] 'x' = : length(v) = 0\n"
## [2] "DEBUG [2015-06-23 20:47:36] 'i': 1\n"
## [3] "DEBUG [2015-06-23 20:47:36] 'i': 0\n"
## [1] "DEBUG [2015-06-23 20:47:36] 'x' = 1,3: length(v) = 2\n"
## [2] "DEBUG [2015-06-23 20:47:36] 'i': 1\n"
## [3] "DEBUG [2015-06-23 20:47:36] 'i': 2\n"
```

```
res
```

```
## [[1]]
## [1] 1.000000 1.732051
##
## [[2]]
## [[2]][[1]]
## [1] NA
##
```

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

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

4.3 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 <- bplapply(list(1, "2", 3), sqrt, BPPARAM = SnowParam(3))
res
```

```
## [[1]]
## [1] 1
##
## [[2]]
## <remote-error in FUN(...): non-numeric argument to mathematical function>
##
## [[3]]
## [1] 1.732051
```

`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 `BPREDO`. The BPPARAM has been changed to `SerialParam` so the job is run locally (in the workspace) and in serial.

```
> bplapply(list(1, "2", 3), fun3, BPREDO = res, 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 problematic input which was the character "2".

5 sessionInfo()

```
toLatex(sessionInfo())
```

```
## \begin{itemize}\raggedright
##   \item R version 3.2.1 (2015-06-18), \verb|x86_64-unknown-linux-gnu|
##   \item Locale: \verb|LC_CTYPE=en_US.UTF-8|, \verb|LC_NUMERIC=C|, \verb|LC_TIME=en_US.UTF-8|, \verb|LC_
##   \item Base packages: base, datasets, grDevices, graphics, methods, parallel,
##     stats, utils
##   \item Other packages: BiocParallel~1.2.6
##   \item Loaded via a namespace (and not attached): BiocStyle~1.6.0,
##     evaluate~0.7, formatR~1.2, futile.logger~1.4.1, futile.options~1.0.0,
##     highr~0.5, knitr~1.10.5, lambda.r~1.1.7, magrittr~1.5, snow~0.3-13,
##     stringi~0.5-2, stringr~1.0.0, tools~3.2.1
## \end{itemize}
```