

Profiling & Parallelization

Lecture 20

Dr. Colin Rundel

Profiling & Benchmarking

profvis demo

```
1 n = 1e6
2 d = tibble(
3   x1 = rt(n, df = 3),
4   x2 = rt(n, df = 3),
5   x3 = rt(n, df = 3),
6   x4 = rt(n, df = 3),
7   x5 = rt(n, df = 3),
8 ) |>
9   mutate(y = -2*x1 - 1*x2 + 0*x3 + 1*x4 + 2*x5 + rnorm(n))
```

```
1 profvis::profvis(lm(y~., data=d))
```

Benchmarking - bench

```
1 d = tibble(  
2   x = runif(10000),  
3   y = runif(10000)  
4 )  
5  
6 (b = bench::mark(  
7   d[d$x > 0.5, ],  
8   d[which(d$x > 0.5), ],  
9   subset(d, x > 0.5),  
10  filter(d, x > 0.5)  
11 ))
```

```
# A tibble: 4 × 6
```

expression	min	median	`itr/sec`	mem_alloc	`gc/sec`
<bch:expr>	<bch:tm>	<bch:tm>	<dbl>	<bch:byt>	<dbl>
1 d[d\$x > 0.5,]	49.4µs	57.4µs	16967.	240.14KB	46.5
2 d[which(d\$x > 0.5),]	90µs	102.6µs	9559.	272.03KB	50.8
3 subset(d, x > 0.5)	87.6µs	104.1µs	9350.	298.36KB	49.3
4 filter(d, x > 0.5)	290.9µs	318µs	3000.	1.48MB	48.7

Larger n

```
1 d = tibble(  
2   x = runif(1e6),  
3   y = runif(1e6)  
4 )  
5  
6 (b = bench::mark(  
7   d[d$x > 0.5, ],  
8   d[which(d$x > 0.5), ],  
9   subset(d, x > 0.5),  
10  filter(d, x > 0.5)  
11 ))
```

```
# A tibble: 4 × 6
```

expression	min	median	`itr/sec`	mem_alloc	`gc/sec`
<bch:expr>	<bch:tm>	<bch:tm>	<dbl>	<bch:byt>	<dbl>
1 d[d\$x > 0.5,]	3.64ms	4.19ms	240.	13.4MB	160.
2 d[which(d\$x > 0.5),]	8.45ms	8.85ms	113.	24.8MB	142.
3 subset(d, x > 0.5)	9.15ms	9.87ms	102.	24.8MB	117.
4 filter(d, x > 0.5)	4.9ms	5.5ms	181.	24.8MB	228.

bench - relative results

```
1 summary(b, relative=TRUE)
```

```
# A tibble: 4 × 6
```

expression	min	median	`itr/sec`	mem_alloc	`gc/sec`
<bch:expr>	<dbl>	<dbl>	<dbl>	<dbl>	<dbl>
1 d[d\$x > 0.5,]	1	1	2.36	1	1.37
2 d[which(d\$x > 0.5),]	2.32	2.11	1.11	1.86	1.21
3 subset(d, x > 0.5)	2.51	2.36	1	1.86	1
4 filter(d, x > 0.5)	1.35	1.31	1.78	1.86	1.95

Parallelization

parallel

Part of the base packages in R

- tools for the forking of R processes (some functions do not work on Windows)
- Core functions:
 - `detectCores`
 - `pvec`
 - `mclapply`
 - `mcpParallel` & `mccollect`

detectCores

Surprisingly, detects the number of cores of the current system.

```
1 detectCores()
```

```
[1] 10
```

pvec

Parallelization of a vectorized function call

```
1 system.time(pvec(1:1e7, sqrt, mc.cores = 1))
```

```
   user  system elapsed  
0.088   0.011   0.099
```

```
1 system.time(pvec(1:1e7, sqrt, mc.cores = 4))
```

```
   user  system elapsed  
0.164   0.118   0.224
```

```
1 system.time(pvec(1:1e7, sqrt, mc.cores = 8))
```

```
   user  system elapsed  
0.091   0.166   0.165
```

```
1 system.time(sqrt(1:1e7))
```

```
   user  system elapsed  
0.017   0.016   0.035
```

pvec - bench::system_time

```
1 bench::system_time(pvec(1:1e7, sqrt, mc.cores = 1))
```

process	real
58.6ms	58.4ms

```
1 bench::system_time(pvec(1:1e7, sqrt, mc.cores = 4))
```

process	real
157ms	199ms

```
1 bench::system_time(pvec(1:1e7, sqrt, mc.cores = 8))
```

process	real
180ms	204ms

```
1 bench::system_time(Sys.sleep(.5))
```

```
process    real  
 60µs     497ms
```

```
1 system.time(Sys.sleep(.5))
```

```
   user  system elapsed  
0.000   0.000   0.505
```

Cores by size

```
1 cores = c(1,4,6,8,10)
2 order = 6:8
3 f = function(x,y) {
4   system.time(
5     pvec(1:(10^y), sqrt, mc.cores = x)
6   ) [3]
7 }
8
9 res = map(
10  cores,
11  function(x) {
12    map_dbl(order, f, x = x)
13  }
14 ) |>
15 do.call(rbind, args = _)
16
17 rownames(res) = paste0(cores, " cores")
18 colnames(res) = paste0("10^",order)
19
20 res
```

	10 ⁶	10 ⁷	10 ⁸
1 cores	0.004	0.057	0.324
4 cores	0.038	0.149	1.738
6 cores	0.031	0.143	1.336
8 cores	0.042	0.137	1.438
10 cores	0.032	0.168	1.406

mclapply

Parallelized version of lapply

```
1 system.time(rnorm(1e7))
```

```
user  system elapsed
0.262  0.004  0.265
```

```
1 system.time(unlist(mclapply(1:10, function(x) rnorm(1e6), mc.cores = 2)))
```

```
user  system elapsed
0.327  0.092  0.268
```

```
1 system.time(unlist(mclapply(1:10, function(x) rnorm(1e6), mc.cores = 4)))
```

```
user  system elapsed
0.335  0.100  0.174
```

```
1 system.time(unlist(mclapply(1:10, function(x) rnorm(1e6), mc.cores = 8)))
```

```
user  system elapsed
0.338  0.150  0.163
```

```
1 system.time(unlist(mclapply(1:10, function(x) rnorm(1e6), mc.cores = 10)))
```

```
user  system elapsed
0.368  0.157  0.169
```

mcpParallel

Asynchronously evaluation of an R expression in a separate process

```
1 m = mcpParallel(rnorm(1e6))  
2 n = mcpParallel(rbeta(1e6,1,1))  
3 o = mcpParallel(rgamma(1e6,1,1))
```

```
1 str(m)
```

List of 2

```
$ pid: int 19229  
$ fd : int [1:2] 4 7  
- attr(*, "class")= chr [1:3] "parallelJob" "childProcess" "process"
```

```
1 str(n)
```

List of 2

```
$ pid: int 19230  
$ fd : int [1:2] 5 9  
- attr(*, "class")= chr [1:3] "parallelJob" "childProcess" "process"
```

mccollect

Checks `mcpaeralel` objects for completion

```
1 str(mccollect(list(m,n,o)))
```

List of 3

```
$ 19229: num [1:1000000] 1.088 0.48 0.706 -2.542 -0.594 ...  
$ 19230: num [1:1000000] 0.192 0.934 0.861 0.64 0.575 ...  
$ 19231: num [1:1000000] 0.25 0.84 1.124 2.366 0.922 ...
```


mccollect - waiting

```
1 p = mcpaerallel(mean(rnorm(1e5)))
```

```
1 mccollect(p, wait = FALSE, 10)
```

```
$`19232`
```

```
[1] 0.0004659567
```

```
1 mccollect(p, wait = FALSE)
```

```
NULL
```

```
1 mccollect(p, wait = FALSE)
```

```
NULL
```

doMC & foreach

doMC & foreach

Packages by Revolution Analytics that provides the `foreach` function which is a parallelizable `for` loop (and then some).

- Core functions:
 - `registerDoMC`
 - `foreach`, `%dopar%`, `%do%`

registerDoMC

Primarily used to set the number of cores used by `foreach`, by default uses `options("cores")` or half the number of cores found by `detectCores` from the `parallel` package.

```
1 options("cores")
```

\$cores

NULL

```
1 detectCores()
```

[1] 10

```
1 getDoParWorkers()
```

[1] 1

```
1 registerDoMC(4)
2 getDoParWorkers()
```

[1] 4

foreach

A slightly more powerful version of base `for` loops (think `for` with an `lapply` flavor). Combined with `%do%` or `%dopar%` for single or multicore execution.

```
1 for(i in 1:10) {  
2   sqrt(i)  
3 }  
4  
5 foreach(i = 1:5) %do% {  
6   sqrt(i)  
7 }
```

```
[[1]]
```

```
[1] 1
```

```
[[2]]
```

```
[1] 1.414214
```

```
[[3]]
```

```
[1] 1.732051
```

[[4]]

[1] 2

foreach - iterators

`foreach` can iterate across more than one value, but it doesn't do length coercion

```
1 foreach(i = 1:5, j = 1:5) %do% {  
2   sqrt(i^2+j^2)  
3 }
```

```
[[1]]  
[1] 1.414214
```

```
[[2]]  
[1] 2.828427
```

```
[[3]]  
[1] 4.242641
```

```
[[4]]  
[1] 5.656854
```

```
[[5]]
```

```
1 foreach(i = 1:5, j = 1:2) %do% {  
2   sqrt(i^2+j^2)  
3 }
```

```
[[1]]  
[1] 1.414214
```

```
[[2]]  
[1] 2.828427
```

foreach - combining results

```
1 foreach(i = 1:5, .combine='c') %do% {  
2   sqrt(i)  
3 }
```

```
[1] 1.000000 1.414214 1.732051 2.000000 2.236068
```

```
1 foreach(i = 1:5, .combine='cbind') %do% {  
2   sqrt(i)  
3 }
```

```
      result.1 result.2 result.3 result.4 result.5  
[1,]          1 1.414214 1.732051          2 2.236068
```

```
1 foreach(i = 1:5, .combine='+') %do% {  
2   sqrt(i)  
3 }
```

```
[1] 8.382332
```


foreach - parallelization

Swapping out `%do%` for `%dopar%` will use the parallel backend.

```
1 registerDoMC(4)
2 system.time(foreach(i = 1:10) %dopar% mean(rnorm(1e6))))
```

```
user  system elapsed
0.298  0.028  0.110
```

```
1 registerDoMC(8)
2 system.time(foreach(i = 1:10) %dopar% mean(rnorm(1e6))))
```

```
user  system elapsed
0.302  0.039  0.078
```

```
1 registerDoMC(10)
2 system.time(foreach(i = 1:10) %dopar% mean(rnorm(1e6))))
```

```
user  system elapsed
0.336  0.051  0.067
```



furrr / future

```
1 system.time( purrr::map(c(1,1,1), Sys.sleep) )
```

```
   user  system elapsed  
0.000   0.000   3.011
```

```
1 system.time( furrr::future_map(c(1,1,1), Sys.sleep) )
```

```
   user  system elapsed  
0.045   0.006   3.074
```

```
1 future::plan(future::multisession) # See also future::multicore  
2 system.time( furrr::future_map(c(1,1,1), Sys.sleep) )
```


```
   user  system elapsed  
0.168   0.004   1.436
```

Example - Bootstrapping

Bootstrapping is a resampling scheme where the original data is repeatedly reconstructed by taking a samples of size n (with replacement) from the original data, and using that to repeat an analysis procedure of interest. Below is an example of fitting a local regression (`loess`) to some synthetic data, we will construct a bootstrap prediction interval for this model.

```
1  set.seed(3212016)
2  d = data.frame(x = 1:120) |>
3    mutate(y = sin(2*pi*x/120) + runif(length(x), -1, 1))
4
5  l = loess(y ~ x, data=d)
6  p = predict(l, se=TRUE)
7
8  d = d |> mutate(
9    pred_y = p$fit,
10   pred_y_se = p$se.fit
11 )
```

```
1 ggplot(d, aes(x,y)) +  
2   geom_point(color="gray50") +  
3   geom_ribbon(  
4     aes(ymin = pred_y - 1.96 * pred_y_se,  
5         ymax = pred_y + 1.96 * pred_y_se),  
6     fill="red", alpha=0.25  
7   ) +  
8   geom_line(aes(y=pred_y)) +  
9   theme_bw()
```



Bootstrapping Demo

What to use when?

Optimal use of parallelization / multiple cores is hard, there isn't one best solution

- Don't underestimate the overhead cost
- Experimentation is key
- Measure it or it didn't happen
- Be aware of the trade off between developer time and run time

BLAS and LAPACK

Statistics and Linear Algebra

An awful lot of statistics is at its core linear algebra.

For example:

- Linear regression models, find

$$\hat{\beta} = (X^T X)^{-1} X^T y$$

- Principle component analysis
 - Find $T = XW$ where W is a matrix whose columns are the eigenvectors of $X^T X$.
 - Often solved via SVD - Let $X = U\Sigma W^T$ then $T = U\Sigma$.

Numerical Linear Algebra

Not unique to Statistics, these are the type of problems that come up across all areas of numerical computing.

- Numerical linear algebra \neq mathematical linear algebra
- Efficiency and stability of numerical algorithms matter
 - Designing and implementing these algorithms is hard
- Don't reinvent the wheel - common core linear algebra tools (well defined API)

BLAS and LAPACK

Low level algorithms for common linear algebra operations

BLAS

- **B**asic **L**inear **A**lgebra **S**ubprograms
- Copying, scaling, multiplying vectors and matrices
- Origins go back to 1979, written in Fortran

LAPACK

- **L**inear **A**lgebra **P**ackage
- Higher level functionality building on BLAS.
- Linear solvers, eigenvalues, and matrix decompositions
- Origins go back to 1992, mostly Fortran (expanded on LINPACK, EISPACK)

Modern variants?

Most default BLAS and LAPACK implementations (like R's defaults) are somewhat dated

- Written in Fortran and designed for a single cpu core
- Certain (potentially non-optimal) hard coded defaults (e.g. block size).

Multithreaded alternatives:


- ATLAS - Automatically Tuned Linear Algebra Software
- OpenBLAS - fork of GotoBLAS from TACC at UTexas
- Intel MKL - Math Kernel Library, part of Intel's commercial compiler tools
- cuBLAS / Magma - GPU libraries from Nvidia and UTK respectively
- Accelerate / vecLib - Apple's framework for GPU and multicore computing

OpenBLAS Matrix Multiply Performance

```
1 x=matrix(runif(5000^2),ncol=5000)
2
3 sizes = c(100,500,1000,2000,3000,4000,5000)
4 cores = c(1,2,4,8,16)
5
6 sapply(
7   cores,
8   function(n_cores)
9     {
10      flexiblas::flexiblas_set_num_threads(n_cores)
11      sapply(
12        sizes,
13        function(s)
14          {
15            y = x[1:s,1:s]
16            system.time(y %**% y)[3]
17          }
18      )
19    }
20 )
```

n	1 core	2 cores	4 cores	8 cores	16 cores
100	0.000	0.000	0.000	0.000	0.000
500	0.004	0.003	0.002	0.002	0.004
1000	0.028	0.016	0.010	0.007	0.009
2000	0.207	0.110	0.058	0.035	0.039
3000	0.679	0.352	0.183	0.103	0.081
4000	1.587	0.816	0.418	0.227	0.145
5000	3.104	1.583	0.807	0.453	0.266

Matrix Multiply of (n x n) matrices



Matrix Multiply of (n x n) matrices

