

Programming with Big Data in R

George Ostrouchov

Oak Ridge National Laboratory and University of Tennessee

Future Trends in Nuclear Physics Computing

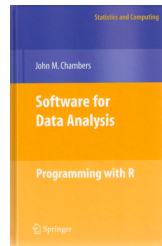
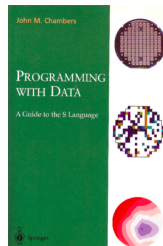
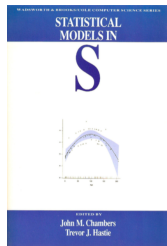
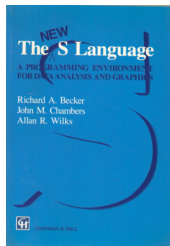
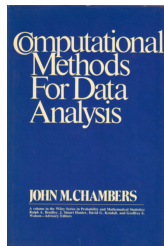
March 16-18, 2016

Thomas Jefferson National Accelerator Facility

Newport News, VA



Why R?: Programming with Data



Chambers.
Computational Methods for Data Analysis. Wiley, 1977.

Becker, Chambers, and Wilks. *The New S Language.* Chapman & Hall, 1988.

Chambers and Hastie. *Statistical Models in S.* Chapman & Hall, 1992.

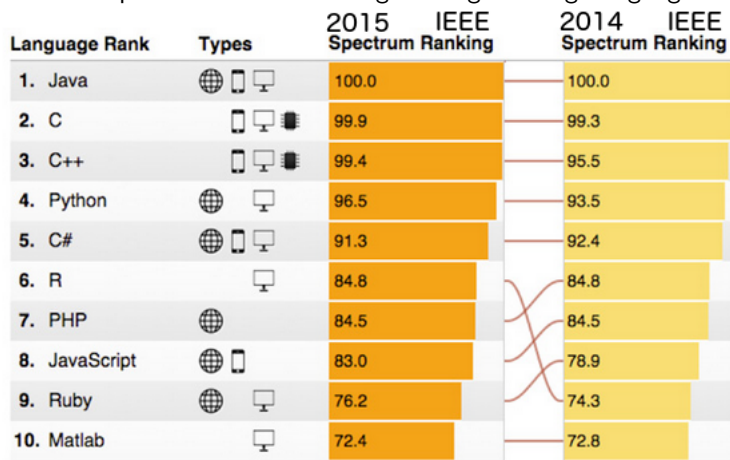
Chambers. *Programming with Data.* Springer, 1998.

Chambers. *Software for Data Analysis: Programming with R.* Springer, 2008.

Thanks to Dirk Eddebuettel for this slide idea and to John Chambers for providing the high-resolution scans of the covers of his books.

Popularity?

IEEE Spectrum's 2014 Ranking of Programming Languages



See: <http://spectrum.ieee.org/static/interactive-the-top-programming-languages#index>

An Example: **knitr** document produced with **RStudio** IDE

Data: 1,653 *start* and *end* timestamps for GPU offload periods.

Want 1-ahead prediction of *start* and *end* to run other codes.

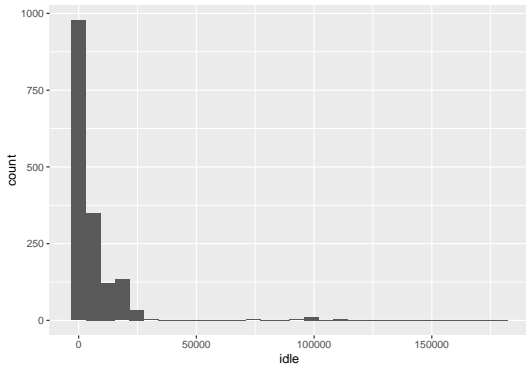
```
ts <- read.table("sortedinode.txt", header=TRUE)
library(dplyr)
ts <- mutate(ts, idle = end - start, busy = start - lag(end))
head(ts)
```

##	start	end	idle	busy
## 1	56114210457	56114211289	832	NA
## 2	56114300920	56114311544	10624	89631
## 3	56114373143	56114373943	800	61599
## 4	56117433146	56117436858	3712	3059203
## 5	56117469818	56117470650	832	32960
## 6	56117507098	56117517081	9983	36448

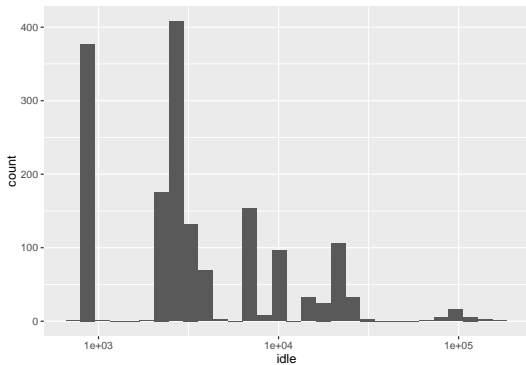
```
dim(ts)
```

```
## [1] 1653    4
```

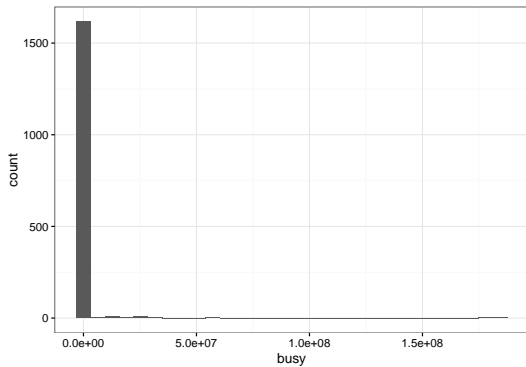
```
library(ggplot2)
ggplot(ts, aes(idle)) + geom_histogram()
```



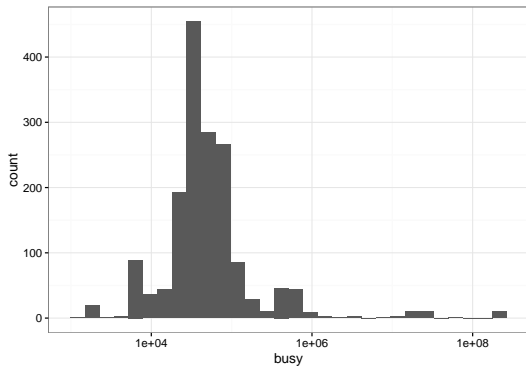
```
ggplot(ts, aes(idle)) + geom_histogram() + scale_x_log10()
```



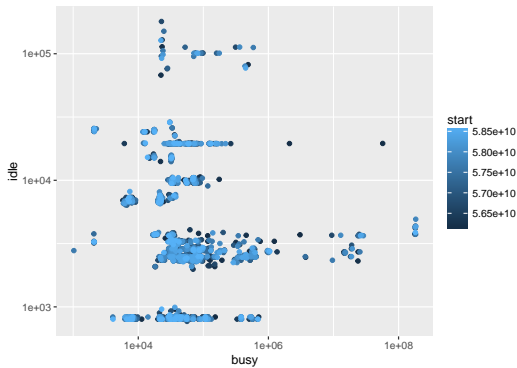
```
ggplot(ts, aes(busy)) + geom_histogram() + theme_bw()
```



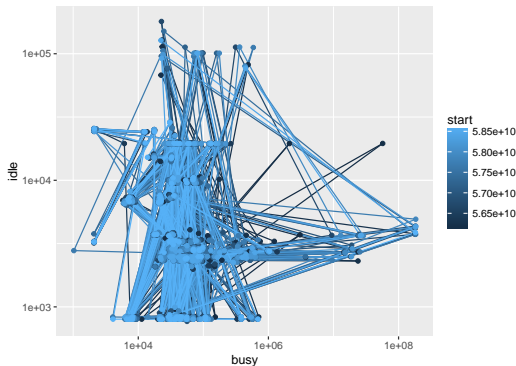
```
ggplot(ts, aes(busy)) + geom_histogram() + scale_x_log10() + theme_bw()
```




```
ggplot(ts, aes(busy, idle, color=start)) + geom_point() + scale_x_log10() +  
  scale_y_log10()
```



```
ggplot(ts, aes(busy, idle, color=start)) + geom_path() + geom_point() +  
  scale_x_log10() + scale_y_log10()
```

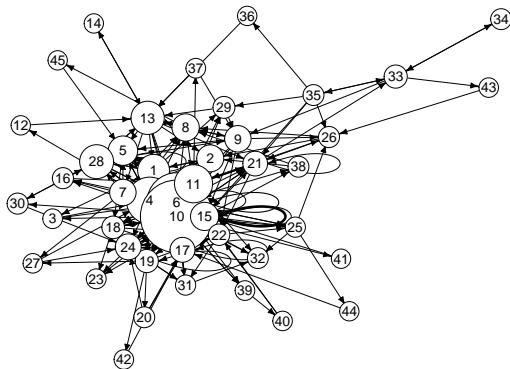


Successive *busy* *idle* periods cluster around several values. Markov chain ... probably sparse ...

Write my own?

Turns out that most of this already exists in the package **rEMM** (One of 7,000+ CRAN packages!)

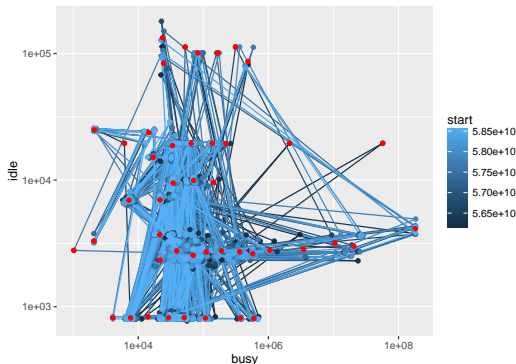
```
ts_log <- transmute(ts, lidle = log10(idle), lbusy = log10(busy))
library(rEMM)
emm <- build(EMM(threshold = 0.2, measure="euclidean"), ts_log)
plot(emm)
```



Markov transition graph after seeing all the data.

Add Markov node locations on top of the log scale *busy-idle* space:

```
ggplot(ts, aes(busy, idle, color=start)) + geom_path() + geom_point() +  
  scale_x_log10() + scale_y_log10() +  
  geom_point(aes(10^lbusy, 10^lidle), data.frame(cluster_centers(emm)), col="red")
```



Looks promising!

```

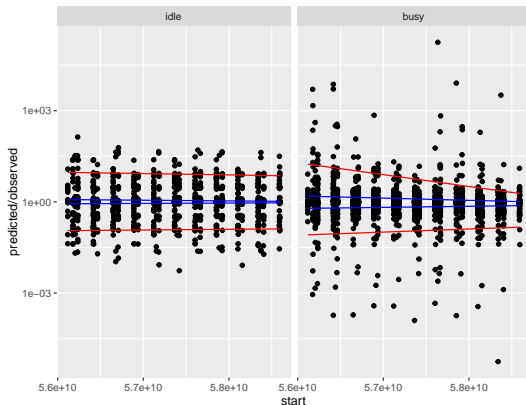
pr_idle <- function(ts, threshold, lambda, measure = "euclidean")
{
  emm <- build(EMM(threshold, measure, lambda = lambda), ts_log[1:5, ])
  ts$idle_p <- ts$busy_p <- NA
  for(i in 6:nrow(ts_log))
  {
    pred_data <- cluster_centers(emm)[predict(emm, n=1), ]
    ts$idle_p[i] <- 10^(pred_data["lidle"])
    ts$busy_p[i] <- 10^(pred_data["lbusy"])
    emm <- build(emm, ts_log[i, ])
  }
  ts
}

ts <- pr_idle(ts, threshold=0.1, lambda=0.01)

```

Prediction proceeds one pair of *busy*, *idle* at a time, predicting the next pair, then updating the Markov states and probabilities for actual values observed.

```
library(reshape2)
ts_melt <- melt(ts, value.name="observed", measure.vars=c("idle", "busy"))
ts_melt <- mutate(ts_melt, predicted=ifelse(variable == "idle", idle_p, busy_p))
ggplot(ts_melt, aes(start, predicted/observed)) + geom_point() + scale_y_log10() +
  stat_quantile(quantiles=c(.75, .25), col="blue") +
  stat_quantile(quantiles=c(.95, .05), col="red") + facet_grid(~variable)
```



This is a lightweight algorithm in R (~2 seconds total for all predictions). It can be made more lightweight (100x ?) by implementing in C or C++ and by updating less often.

Resources for Learning R

- RStudio IDE
<http://www.rstudio.com/products/rstudio-desktop/>
- Task Views: <http://cran.at.r-project.org/web/views>
- Book: *The Art of R Programming* by Norm Matloff:
<http://nostarch.com/artofr.htm>
- *Advanced R*: <http://adv-r.had.co.nz/> and *ggplot2*
<http://docs.ggplot2.org/current/> by Hadley Wickham
- R programming for those coming from other languages: http://www.johndcook.com/R_language_for_programmers.html
- *aRrgh: a newcomer's (angry) guide to R*, by Tim Smith and Kevin Ushey: <http://tim-smith.us/arrgh/>
- Mailing list archives: <http://tolstoy.newcastle.edu.au/R/>
- The [R] stackoverflow tag.

Why R?:

Programming with **Big** Data

Wei-Chen Chen¹

George Ostrouchov^{2,3,4}

Pragneshkumar Patel³

Drew Schmidt⁴

¹FDA
Washington, DC, USA

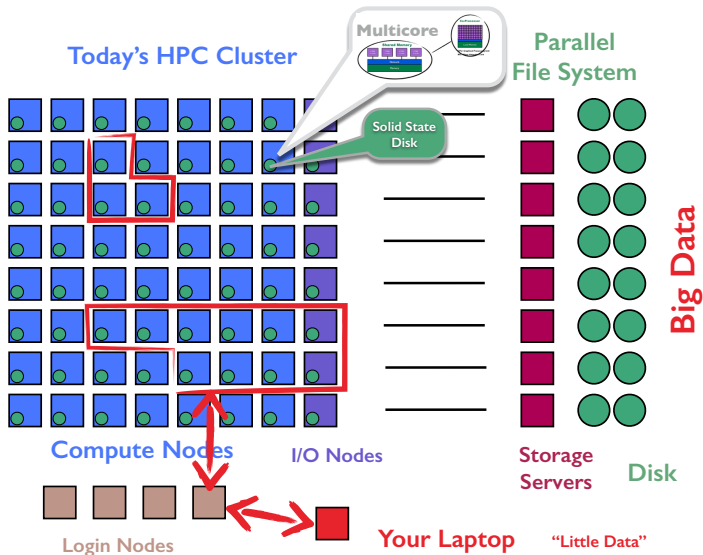
²Computer Science and Mathematics Division
Oak Ridge National Laboratory, Oak Ridge TN, USA



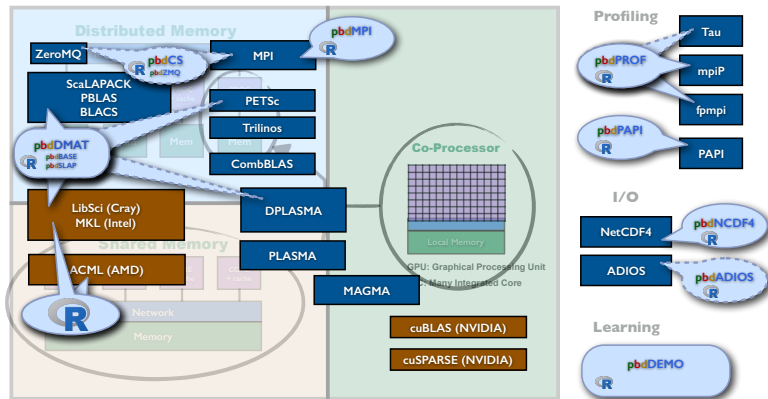
³Joint Institute for Computational Sciences
University of Tennessee, Knoxville TN, USA

⁴Business Analytics and Statistics
University of Tennessee, Knoxville TN, USA

HPC Cluster with NVRAM and Parallel File System



pbdR Interfaces to Libraries: Sustainable Path



Why use HPC libraries?

- Many science communities are invested in their API.
- Data analysis uses much of the same basic math as simulation science
- The libraries represent 30+ years of parallel algorithm research
- *They're tested. They're fast. They're scalable.*

pbdMPI: a High Level Interface to MPI

- API is simplified: defaults in control objects.
- S4 methods: extensible to complex R objects.
- Additional error checking
- Array and matrix methods without serialization: faster than **Rmpi**.

pbdMPI (S4)	Rmpi
<code>allreduce</code>	<code>mpi.allreduce</code>
<code>allgather</code>	<code>mpi.allgather</code> , <code>mpi.allgatherv</code> , <code>mpi.allgather.Robj</code>
<code>bcast</code>	<code>mpi.bcast</code> , <code>mpi.bcast.Robj</code>
<code>gather</code>	<code>mpi.gather</code> , <code>mpi.gatherv</code> , <code>mpi.gather.Robj</code>
<code>recv</code>	<code>mpi.recv</code> , <code>mpi.recv.Robj</code>
<code>reduce</code>	<code>mpi.reduce</code>
<code>scatter</code>	<code>mpi.scatter</code> , <code>mpi.scatterv</code> , <code>mpi.scatter.Robj</code>
<code>send</code>	<code>mpi.send</code> , <code>mpi.send.Robj</code>

SPMD: Copies of One Code Run Asynchronously

A simple SPMD allreduce

allreduce.r

```
1 library(pbdMPI, quiet = TRUE)
2
3 ## Your local computation
4 n <- comm.rank() + 1
5
6 ## Now "Reduce" and give the result to all
7 all_sum <- allreduce(n) # Sum is default
8
9 text <- paste("Hello: n is", n, "sum is", all_sum )
10 comm.print(text, all.rank=TRUE)
11
12 finalize()
```

Execute this batch script via:

```
1 mpirun -np 2 Rscript allreduce.r
```

Output:

```
1 COMM.RANK = 0
2 [1] "Hello: n is 1 sum is 3"
3 COMM.RANK = 1
4 [1] "Hello: n is 2 sum is 3"
```

Machine Learning Example: Random Forest

Example: Letter Recognition data from package **mlbench** ($20,000 \times 17$)



1	[,1]	lettr	capital letter
2	[,2]	x.box	horizontal position of box
3	[,3]	y.box	vertical position of box
4	[,4]	width	width of box
5	[,5]	high	height of box
6	[,6]	onpix	total number of on pixels
7	[,7]	x.bar	mean x of on pixels in box
8	[,8]	y.bar	mean y of on pixels in box
9	[,9]	x2bar	mean x variance
10	[,10]	y2bar	mean y variance
11	[,11]	xybar	mean x y correlation
12	[,12]	x2ybr	mean of x^2 y
13	[,13]	xy2br	mean of x y^2
14	[,14]	x.ege	mean edge count left to right
15	[,15]	xegvy	correlation of x.ege with y
16	[,16]	y.ege	mean edge count bottom to top
17	[,17]	yegvx	correlation of y.ege with x

P. W. Frey and D. J. Slate (Machine Learning Vol 6/2 March 91): "Letter Recognition Using Holland-style Adaptive Classifiers".

Example: Random Forest Code (build many simple models, use model averaging to predict)

Serial Code 4_rf_s.r

```
1 library(randomForest)
2 library(mlbench)
3 data(LetterRecognition) # 26 Capital Letters Data 20,000 x 17
4 set.seed(seed=123)
5 n <- nrow(LetterRecognition)
6 n_test <- floor(0.2*n)
7 i_test <- sample.int(n, n_test) # Use 1/5 of the data to test
8 train <- LetterRecognition[-i_test, ]
9 test <- LetterRecognition[i_test, ]
10
11 ## train random forest
12 rf.all <- randomForest(letter ~ ., train, ntree=500,
13   norm.votes=FALSE)
14
15 ## predict test data
16 pred <- predict(rf.all, test)
17 correct <- sum(pred == test$letter)
18 cat("Proportion Correct:", correct/(n_test), "\n")
```

Example: Random Forest Code

(Split learning by blocks of trees. Split prediction by blocks of rows.)

Parallel Code 4_rf_p.r

```
1 library(randomForest)
2 library(mlbench)
3 data(LetterRecognition)
4 comm.set.seed(seed=123, diff=FALSE) # same training data
5 n <- nrow(LetterRecognition)
6 n_test <- floor(0.2*n)
7 i_test <- sample.int(n, n_test) # Use 1/5 of the data to test
8 train <- LetterRecognition[-i_test, ]
9 test <- LetterRecognition[i_test, ][get.jid(n_test), ]
10
11 comm.set.seed(seed=1e6*runif(1), diff=TRUE)
12 my.rf <- randomForest(lettr ~ ., train, ntree=500%%comm.size(),
13   norm.votes=FALSE)
14
15 rf.all <- do.call(combine, allgather(my.rf))
16
17 pred <- predict(rf.all, test)
18 correct <- allreduce(sum(pred == test$lettr))
19 comm.cat("Proportion Correct:", correct/(n_test), "\n")
```

Dense Matrix and Vector Operations

A matrix is mapped to a processor grid shape

$$\begin{bmatrix} 0 & 1 & 2 & 3 & 4 & 5 \end{bmatrix}$$

(a) 1×6

$$\begin{bmatrix} 0 & 1 & 2 \\ 3 & 4 & 5 \end{bmatrix}$$

(b) 2×3

$$\begin{bmatrix} 0 & 1 \\ 2 & 3 \\ 4 & 5 \end{bmatrix}$$

(c) 3×2

$$\begin{bmatrix} 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{bmatrix}$$

(d) 6×1

Table: Processor Grid Shapes with 6 Processors

pbdR No change in syntax.

Data redistribution functions.

```
1 x <- x[-1, 2:5]
2 x <- log(abs(x) + 1)
3 x.pca <- prcomp(x)
4 xtx <- t(x) %*% x
5 ans <- svd(solve(xtx))
```

*The above (and over 100 other functions) runs on 1 core with R
or 10,000 cores with **pbdR** ddmatrix class*

```
1 > showClass("ddmatrix")
2 Class "ddmatrix" [package "pbdDMAT"]
3 Slots:
4 Name:      Data      dim      ldim      bldim      ICTXT
5 Class:     matrix numeric numeric numeric numeric
```

```
1 > x <- as.rowblock(x)
2 > x <- as.colblock(x)
3 > x <- redistribute(x, bldim=c(8, 8), ICTXT = 0)
```

Truncated SVD from random projections¹

PROTOTYPE FOR RANDOMIZED SVD

Given an $m \times n$ matrix A , a target number k of singular vectors, and an exponent q (say, $q = 1$ or $q = 2$), this procedure computes an approximate rank- $2k$ factorization $U\Sigma V^*$, where U and V are orthonormal, and Σ is nonnegative and diagonal.

Stage A:

- 1 Generate an $n \times 2k$ Gaussian test matrix Ω .
- 2 Form $Y = (AA^*)^q A\Omega$ by multiplying alternately with A and A^* .
- 3 Construct a matrix Q whose columns form an orthonormal basis for the range of Y .

Stage B:

- 4 Form $B = Q^* A$.
- 5 Compute an SVD of the small matrix: $B = \tilde{U}\tilde{\Sigma}V^*$.
- 6 Set $U = Q\tilde{U}$.

Note: The computation of Y in step 2 is vulnerable to round-off errors. When high accuracy is required, we must incorporate an orthonormalization step between each application of A and A^* ; see Algorithm 4.4.

ALGORITHM 4.4: RANDOMIZED SUBSPACE ITERATION

Given an $m \times n$ matrix A and integers ℓ and q , this algorithm computes an $m \times \ell$ orthonormal matrix Q whose range approximates the range of A .

- 1 Draw an $n \times \ell$ standard Gaussian matrix Ω .
- 2 Form $Y_0 = A\Omega$ and compute its QR factorization $Y_0 = Q_0 R_0$.
- 3 **for** $j = 1, 2, \dots, q$
- 4 Form $\tilde{Y}_j = A^* Q_{j-1}$ and compute its QR factorization $\tilde{Y}_j = \tilde{Q}_j \tilde{R}_j$.
- 5 Form $Y_j = A Q_j$ and compute its QR factorization $Y_j = Q_j R_j$.
- 6 **end**
- 7 $Q = Q_q$.

Serial R

```

1 randSVD <- function(A, k, q=3)
2 {
3   ## Stage A
4   Omega <- matrix(rnorm(n*2*k),
5     nrow=n, ncol=2*k)
6   Y <- A %%% Omega
7   Q <- qr.Q(qr(Y))
8   At <- t(A)
9   for(i in 1:q)
10    {
11      Y <- At %%% Q
12      Q <- qr.Q(qr(Y))
13      Y <- A %%% Q
14      Q <- qr.Q(qr(Y))
15    }
16
17   ## Stage B
18   B <- t(Q) %%% A
19   U <- La.svd(B)$u
20   U <- Q %%% U
21   U[, 1:k]
22 }
```

¹Halko, Martinsson, and Tropp. 2011. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions *SIAM Review* 53 217–288

Truncated SVD from random projections

Serial R

```

1 randSVD <- function(A, k, q=3)
2 {
3   ## Stage A
4   Omega <- matrix(rnorm(n*2*k),
5     nrow=n, ncol=2*k)
6   Y <- A %*% Omega
7   Q <- qr.Q(qr(Y))
8   At <- t(A)
9   for(i in 1:q)
10     {
11       Y <- At %*% Q
12       Q <- qr.Q(qr(Y))
13       Y <- A %*% Q
14       Q <- qr.Q(qr(Y))
15     }
16   ## Stage B
17   B <- t(Q) %*% A
18   U <- La.svd(B)$u
19   U <- Q %*% U
20   U[, 1:k]
21 }

```

Parallel pbdR

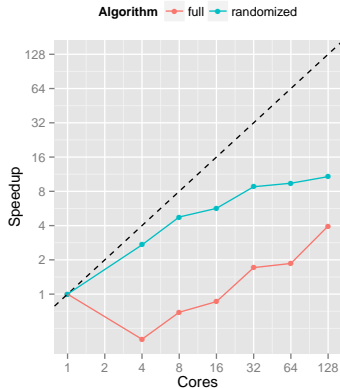
```

1 randSVD <- function(A, k, q=3)
2 {
3   ## Stage A
4   Omega <- ddmatrix("rnorm",
5     nrow=n, ncol=2*k)
6   Y <- A %*% Omega
7   Q <- qr.Q(qr(Y))
8   At <- t(A)
9   for(i in 1:q)
10     {
11       Y <- At %*% Q
12       Q <- qr.Q(qr(Y))
13       Y <- A %*% Q
14       Q <- qr.Q(qr(Y))
15     }
16   ## Stage B
17   B <- t(Q) %*% A
18   U <- La.svd(B)$u
19   U <- Q %*% U
20   U[, 1:k]
21 }

```

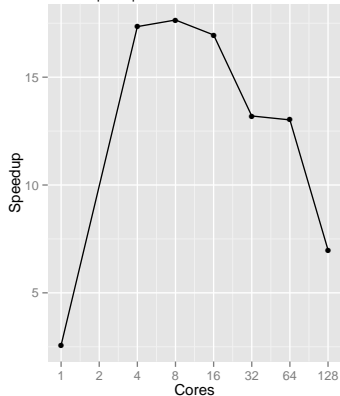
From journal to scalable code and scaling data in one day.

30 Singular Vectors from a 100,000 by 1,000 Matrix



Speedup relative to 1 core

30 Singular Vectors from a 100,000 by 1,000 Matrix
Speedup of Randomized vs. Full SVD



RandSVD speedup relative to full SVD

Future Work

- NSF/DMS: second year of a 3 year grant to
 - Bring back interactivity via client/server (pbdCS/pbdZMQ)
 - Simplify parallel data input
 - Begin DPLASMA integration
 - Outreach to the statistics community
- DOE/SC: In-situ or staging use with simulations
 - pbdADIOS - HPC I/O
- Pending: NSF BIGDATA, Tensor Regression
- Pending: Exascale Computing Project, analytics for ParaView/VisIt

Where to learn more?

- <http://r-pbd.org/>
- **pbdDEMO** vignette
- [Googlegroup:RBigDataProgramming](#)
- **pbdR** Installations: OLCF, NERSC, SDSC, TACC, IU, BSC Spain, CSCS Switzerland, IT4I Czech, ISM Japan, and many more
- Need access to a cluster computer? From NSF:
 - XSEDE *trial* or *startup* allocation
<https://www.xsede.org/web/xup/allocations-overview>.
 - Most resources have **pbdR** installed

Support

This material is based upon work supported by the National Science Foundation Division of Mathematical Sciences under Grant No. 1418195.

This work used resources of the [Oak Ridge Leadership Computing Facility](#) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

This work also used resources of [National Institute for Computational Sciences](#) at the University of Tennessee, Knoxville, which is supported by the U.S. National Science Foundation.

