Big Data Analytics and HPC

Matthew J. Denny

mzd5530@psu.edu – www.mjdenny.com – @MatthewJDenny

www.mjdenny.com/ICPSR_Data_Science_2015.html

July 28, 2015

INSTITUTE FOR SOCIAL SCIENCE RESEARCH
UNIVERSITY OF MASSACHUSETTS AMHERST

PENNSTATE DEPARTMENT OF POLITICAL SCIENCE

BDSS

Overview

1. Overview of High Performance Computing/Big Data Analytics in R.
2. Programming Choices.
3. Big Data Example.
4. Hardware.

High Performance Computing

- Make use of low overhead, high speed programming languages (`C, C++, Fortran, Java, etc.`)
- Parallelization
- Efficient implementation.
- **Good scheduling**.

Big Data Analytics

- Use memory efficient data structures and programming languages.
- More RAM.
- Databases.
- Efficient inference procedures.
- **Good scheduling**.

How They Fit Together

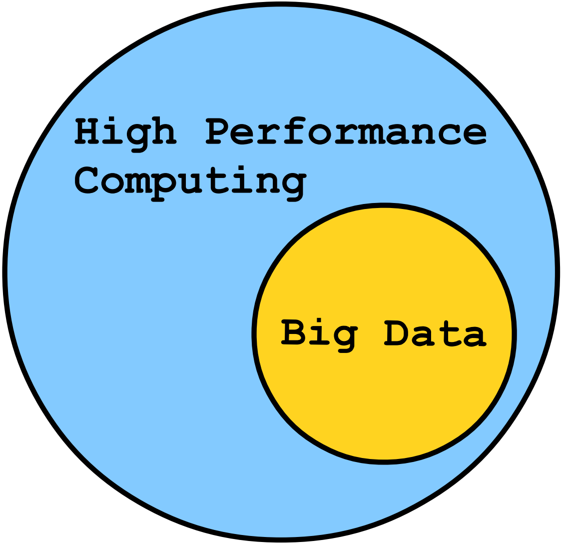

Hardware Constraints

- **RAM** = computer working memory – determines size of datasets you can work on.

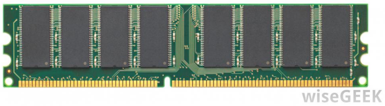

- **CPU** = processor, determines speed of analysis and degree of parallelization.

Look At Your Activity Monitor!

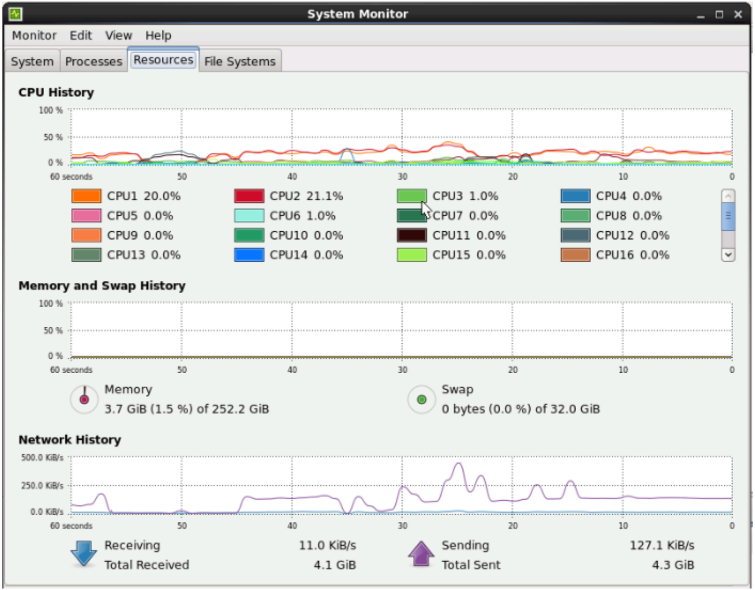

2. Programming Choices

Efficient R Programming

- Loops are slow in R, but fast enough for most things.
- Built-in functions are mostly written in C – much faster!
- Subset data before processing when possible.
- Avoid growing data structures.

Loops Are “Slow” in R

```
system.time({
    vect <- c(1:10000000)
    total <- 0
    #check using a loop
    for(i in 1:length(as.numeric(vect))){
        total <- total + vect[i]
    }
    print(total)
})
[1] 5e+13
   user  system elapsed
  7.641   0.062   7.701
```

And Dast in C

```
system.time({
    vect <- c(1:10000000)
    #use the builtin R function
    total <- sum(as.numeric(vect))
    print(total)
})
[1] 5e+13
   user  system elapsed
  0.108   0.028   0.136
```

Summing Over A Sparse Dataset

```
#number of observations
numobs <- 100000000

#observations we want to check
vec <- rep(0,numobs)

#only select 100 to check
vec[sample(1:numobs,100)] <- 1

#combine data
data <- cbind(c(1:numobs),vec)
```

Conditional Checking

```
system.time({
    total <- 0
    for(i in 1:numobs){
        if(data[i,2] == 1)
        total <- total + data[i,1]
    }
    print(total)
})
[1] 5385484508
   user  system elapsed
199.917   0.289 200.350
```

Subsetting

```
system.time({
    dat <- subset(data, data[,2] ==1)
    total <- sum(dat[,1])
    print(total)
})
[1] 5385484508
   user  system elapsed
  5.474   1.497   8.245
```

2.a. Pre-Allocation

Adding To A Vector vs. Pre-Allocation

```
system.time({
  vec <- NULL
  for (i in 1:(10^5)) vec <- c(vec,i)
})
   user  system elapsed
 18.495   7.401  25.935
system.time({
  vec <- rep(NA,10^5)
  for (i in 1:(10^5)) vec[i] <- i
})
   user  system elapsed
  0.144   0.002   0.145
```

Pre-Allocated Vector – Bigger Example

```
system.time({
  vec <- rep(NA,10^6)
  for (i in 1:(10^6)) vec[i] <- i
})
   user  system elapsed
  1.765   0.040   1.872
```

Adding To A Vector – Bigger Example

```
system.time({
  vec <- NULL
  for (i in 1:(10^6)) vec <- c(vec,i)
})
```

```
Timing stopped at: 924.922 120.322 1872.294
```

I didn't feel like waiting...

Pre-Allocation

- Vectors in R can only hold about 2.1 Billion elements.
- Write to over-allocated vector then subset.
- Speedup is exponential in the vector size and number of additions.

2.b. Parallelization

Parallelization Using `foreach`

- Works best when we need to calculate some complex statistic on each row/column of dataset.
- Works just like a regular `for( )` loop as long as operations are **independent**.
- Good for bootstrapping.

Parallelization Using foreach

```
# Packages:
require(doMC)
require(foreach)

# Register number of cores
nCores <- 8
registerDoMC(nCores)

# iterations
N <- 100

# Run analysis in parallel
results <- foreach(i=1:N,.combine=rbind) %dopar% {
    result <- function(i)
}
```

Parallelization Using A `snowfall` Cluster

- Can run across multiple machines.
- Can run totally different jobs on each thread.
- Requires explicit management by researcher.

Parallelization Using A snowfall Cluster

```
# Package:
library(snowfall)

# Register cores
numcpus <- 4
sfInit(parallel=TRUE, cpus=numcpus )

# Check initialization
if(sfParallel()){
    cat( "Parallel on", sfCpus(), "nodes.\n" )
}else{
    cat( "Sequential mode.\n" )
}
```

Parallelization Using A `snowfall` Cluster

```
# Export all packages
for (i in 1:length(.packages())){
     eval(call("sfLibrary", (.packages()[i]),
     character.only=TRUE))
}

# Export a list of R data objects
sfExport("Object1","Object2","Object3")

# Apply a function across the cluster
result <- sfClusterApplyLB(indexes,Function)

# Stop the cluster
sfStop()
```

Parallelization Using `mclapply( )`

- **Will not work with Windows machines.**
- Simple parallelization.
- Works well with functions written in Rcpp.

Parallelization Using `mclapply( )`

```
# Packages:
library(parallel)

# Wrapper Function
run_on_cluster <- function(i){
    temp <- your_function(i,some other stuff)
    return(temp)
}

# Run analysis
indexes <- 1:Iterations
Result  <- mclapply(indexes,
                    run_on_cluster,
                    mc.cores = num_cpus)
```

2.c. Memory Efficient Regression

High Memory Regression Using `biglm( )`

- Memory efficient implementation of `glm( )`
- Can also read in data in chunks from internet or a relational database.
- Will not take data in matrix form, only data.frame

High Memory Regression Using `biglm( )`

```
# Data must be of data.frame type
data <- as.data.frame(data)

# Use variable names in formula
str <- "V1 ~ V2 + V4"

# Run model
model<- bigglm(as.formula(str),
               data = full_data,
               family=binomial(),
               maxit = 20)
```

3. Big Data Example

Congressional Bill Text

1. 90,867 final versions of bills introduced in the U.S. Congress form 1993-2012.
2. 293,697,605 tokens (370,445 unique).
3. 90 covariates for every bill.
4. Addition data on amendments, cosponsorships, and floor speeches.

Lets Look at Some Bill Text

```
    This Act may be cited as the ``EPA Science Act of 2014''.
    .....
SEC. 2. SCIENCE ADVISORY BOARD.

   (a) Independent Advice.--Section 8(a) of the Environmental
Research, Development, and Demonstration Authorization Act
of 1978 (42 U.S.C. 4365(a)) is amended by inserting
``independently'' after ``Advisory Board which shall''.
   (b) Membership.--Section 8(b) of the Environmental Research,
Development, and Demonstration Authorization Act of 1978
(42 U.S.C. 4365(b)) is amended to read as follows:
        ``(b)(1) The Board shall be composed of at least nine
        members, one of whom shall be designated Chairman, and
        shall meet at such times and places as may be designated
        by the Chairman.
       ``(2) Each member of the Board shall be qualified
       by education, training, and experience to evaluate
       scientific and technical ....
```

Taxonomy of Bill Text

Category	Definition	Example
Substantive Language	Confers the intent of a piece of legislation or a particular provision.	{**restrict abortion**}, {**reduce the deficit**}
Topical Language	Confers information about the subject of the Bill.	{alternate academic achievement standards}
Boilerplate	Gives direction about legal interpretation or implementation.	{Notwithstanding any other provision of this paragraph...}
Domain Stopwords	Gives no information about intent, legal interpretation or implementation.	{SECTION}, {(c) Title III.–}, {(1) Part a.–}, {(A) Subpart 1.–}, {to adopt}

Lets Look at N-Grams

Unit of Analysis	Topical Text
Tokens	{Should}, {a}, {Federal}, {agency}, {seek}, {to}, {**restrict**}, {**photography**}, {of}, {its}, { installations}, {or}, {personnel}, {it}, {shall}, {obtain}, {a}, {**court**}, {order}, {that}, {outlines}, { the}, {national}, {**security**}, {or}, {other}, {reasons}, {for}, {the}, {restriction}
Bigrams	{Should a}, {a Federal}, {Federal agency}, {agency seek}, {seek to}, {to restrict}, {**restrict photography**}, {photography of}, {of its}, {its installations}, {installations or}, {or personnel}, {personnel it}, {it shall}, {shall obtain}, {obtain a}, {a court}, {**court order**}, {order that}, {that outlines}, {outlines the}, {the national}, {**national security**}, {security or}, {or other}, {other reasons}, {reasons for}, {for the}, {the restriction}

Syntactic Filtering

Tag Pattern	Example
AN	*linear function*
NN	*regression coefficients*
AAN	*Gaussian random variable*
ANN	*cumulative distribution function*
NAN	*mean squared error*
NNN	*class probability function*
NPN	*degrees of freedom*

Verbs capture actions...

Tag Pattern	Example
VN	*reduce funding*
VAN	*encourage dissenting members*
VNN	*restrict federal agencies*

Syntactic Filtering and Phrase Extraction

Unit of Analysis	Matches
Filtered Bigrams	{Federal agency}, {**restrict photography**}, {**court order**}, {**national security**}, {other reasons}
Filtered Trigrams	NONE
Noun Phrases	{Federal agency}, {**court order**}, {**national security**},{other reasons},{**other reasons for the restriction**}

Constructing a Document-Term Matrix

- Want Document x Vocabulary matrix.
- Take advantage of sparsity.
- Use C++ for indexing.
- Have to chunk and add.

Constructing A Document-Term Matrix

```
# gives us simple triplet matrix class
library(slam)

# load in C++ function to generate rows in matrix
Rcpp::sourceCpp('Document_Word_Compiler.cpp')

for(i in 1:chunks){
  dwm <- Generate_Document_Word_Matrix(chunk,..)
  dwm <- as.simple_triplet_matrix(dwm)
  if(j == 1){
    swm <- dwn
  }else{
    swm <- rbind(swm,dwm)
  }
}
```

Semi-Supervised Major Topic Tags

Category
1. Macroeconomics
2. Civil Rights, Minority Issues, and Civil Liberties
3. Health
4. Agriculture
5. Labor and Employment
6. Education
7. Environment
8. Energy
9. Immigration
10. Transportation
12. Law, Crime, and Family Issues
13. Social Welfare
14. Community Development and Housing Issues
15. Banking, Finance, and Domestic Commerce
16. Defense
.....

Filtered Trigrams for Identifying Topical Area

Health

Word	**PMI**	**Local**	**Global**
stay	1.754	24425	26428
start	1.689	12395	14321
enhance	1.684	22142	25707
providers	1.684	51656	59976
mining	1.679	15221	17751

Filtered Trigram	**PMI**	**Local**	**Global**
health insurance plan	1.340	868	1528
health benefit plan	1.306	1125	2049
term care insurance	1.292	3564	6577
health plan means	1.110	261	578
alternative dispute resolution	1.079	923	2108

Big Data Challenges

1. Extending R vectors/matrices beyond 2.1 billion elements.
2. More low level data structures – linked list, queue, stack, etc.
3. Better garbage collection.
4. More lazy/greedy/approximate methods.

4. Hardware

Classes of Hardware

1. Supercomputers
2. Mainframes
3. **Cluster Computing Resources**
4. Servers
5. **HPC Workstations**
6. **Consumer Desktops**
7. GPGPU

Supercomputers

- Used when all computing resources are needed to solve one problem.
- Physics, engineering, materials science

Mainframes

- Used for large database applications.
- Business analytics, healthcare.

Cluster Computing Resources

- Flexible, used for parallel and high memory tasks.
- General purpose academic computing infrastructure.

Servers

- Most often used for hosting websites.
- Can be useful for long jobs, high memory.

HPC Workstations

- Personal mid-size high memory and parallel computing.
- For people who moderate resources **constantly**.

Desktop

- Everything, depending on how long you are willing to wait.
- Will run 95% of what you want to do.

General Purpose GPU Computing

- Problems that break down to small, interdependent parts.
- Bootstrapping, complex looping, optimization.

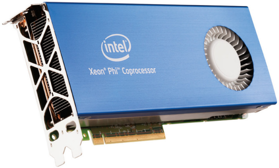

Pricing Tiers

- **Cluster Access :** Usually free through your institution but often requires application/faculty sponsorship.
- **HPC Workstation:** \$8,000-\$15,000 – Not a good investment for most.
- **Desktop:** \$700-\$2000 – Often a very good investment.

My Suggestion

- Ask a faculty member for access. **TRY THIS FIRST.**
- Investing in a desktop with 4C/8T (Intel i7) and 16GB+ of RAM is often a smart idea if it will not get in the way of conference attendance.
- Access a university cluster only once you are confident a desktop can no longer meet your needs.

Upgrades

- Old computers are good for HPC tasks that simply take a while to run.
- Locate computer in academic office for free electricity/internet/easier remote access.
- Relatively cheap upgrades can dramatically improve performance.
- **BENCHMARK**

Know Your Motherboard

- How many RAM Slots?
- Peripherals, CPU, GPU slots.

RAM

- Work with larger datasets.
- 16GB Kits – (\$100-150)
- 32GB Kits – (\$200-300)

Solid State Hard Drive

- General system performance, data I/O.
- \$0.40-\$0.80 per GB
- Leave 15-20% free.

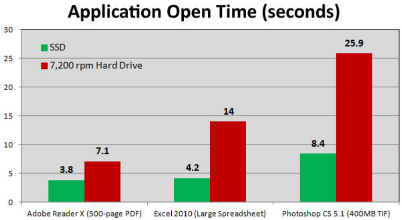

Traditional hard disk drive

Solid state hard drive

Check Review Sites

ANANDTECH

tom's hardware

Things To Remember

- Get an Uninterrupted Power Supply (UPS) for stable power.

- Put a sign on your computer that says don't touch.

Summary

- Don't buy it unless you **absolutely need it**.
- Most resources can be borrowed/had for free.
- More powerful resources require **more time** to learn.