

Parallel External Memory Algorithms in RevoPemaR and RevoScaleR

Presented by:

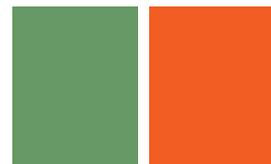
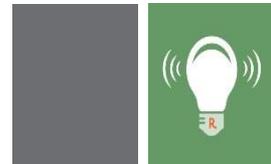
Mario Inchiosa, Ph.D.

Chief Scientist, Revolution Analytics

Workshop on Distributed Computing in R

HP Labs, Palo Alto, CA

January 27, 2015





Agenda

- Introduction to PEMAs
- RevoPemaR Package
- RevoScaleR Package



Introduction

- In many fields, the **size of data sets** is increasing more rapidly than the speed of single cores, of RAM, and of hard drives
- In addition, data is being collected and stored in **distributed locations**
- In some applications it is also important to be able to **update computations as new data is obtained**
- To deal with this, there is increasing need for statistical and machine learning code that does not require all data to be in memory and that can **distribute computations across cores, computers, and time**
- **Parallel External Memory Algorithms** provide a foundation for such software



External memory algorithms (EMA's)

- Algorithms that do not require all data to be in RAM
- Data is processed sequentially, a chunk at a time, with intermediate results produced for each chunk
- Intermediate results for one chunk of data can be combined with those of another chunk
- Such algorithms are widely available for statistical and machine learning methods
 - Regression, Classification
 - Unsupervised Learning
 - Transformation



Example external memory algorithm for the mean of a variable

- **initialize** method: $\text{sum}=0$, $\text{totalObs}=0$
- **processData** method: for each chunk of x ; $\text{sum} = \text{sum} + \text{sum}(x)$,
 $\text{totalObs} = \text{totalObs} + \text{length}(x)$
- **updateResults** method:
 - $\text{sum}_{12} = \text{sum}_1 + \text{sum}_2$
 - $\text{totalObs}_{12} = \text{totalObs}_1 + \text{totalObs}_2$
- **processResults** method: $\text{mean} = \text{sum} / \text{totalObs}$



Parallel external memory algorithms (PEMA's)

- PEMA's are implementations of EMA's that allow them to be executed on multiple cores and computers
- Chunks of data are processed in parallel, as well as sequentially
- To write such algorithms, the inherently sequential parts of the computation must be separated from the parallelizable parts
- For example, for iterative algorithms such a maximum likelihood (e.g. IRLS for glm), each iteration depends upon the previous one so the iterations cannot be run in parallel. However, the computations for each iteration can often be parallelized, and this is usually where most of the time is spent in any case.
- Keys to efficiency
 - avoid inter-thread and inter-process communication as much as possible
 - rapidly get chunks of data to the processData methods



Reference classes

- Object oriented system in R that is similar to those in Java, C++
- Fields of RC objects are mutable; they are not copied upon modification
- Methods belong to objects, can operate on the fields of the object, and all methods see the changes made by any method to a field
- Introduced in R 2.12, by John Chambers
- Is a special S4 class that wraps an environment



Why reference classes?

- Efficiency and lowered memory usage
 - RC's allow control over when objects are copied
 - Member variables (fields) do not have to be passed as arguments to other methods
- Encapsulation
 - The data and the methods that operate on them are bundled
 - Access methods and field locking provide control over access to fields
- Convenience
 - It is possible to do the same things in other ways in R, but RC's are more convenient
- Familiarity
 - Familiar OOP system for Java, C++, and other programmers

RevoPemaR



- An R package providing a framework for writing Parallel External Memory Algorithms (PEMA's)
- Uses Reference Classes
- Code can be written, tested, and run using data in memory on one computer, and then can be deployed to a variety of distributed compute contexts, using a variety of data sources
- Versions of this package are available as part of Revolution R Enterprise 7.2 and 7.3, under the Apache 2.0 license
- This framework is based on, and is very similar to, a C++ framework we have used for several years in RevoScaleR to implement extremely high performance statistical and machine learning algorithms



Features of code written using the RevoPemaR framework

- **Scalable:** it can process an unlimited number of rows of data in a fixed amount of RAM, even on a single core
- **Distributable:** it is distributable across processes on a single computer and across nodes of a cluster
- **Updateable:** existing results can be updated given new data
- **Portable:**
 - **With respect to platforms:** it can be executed on a wide variety of computing platforms, including the parallel and distributed platforms supported by Revolution's RevoScaleR package (Teradata, IBM Platform LSF, Microsoft HPC Server, and various flavors of Hadoop)
 - **With respect to data sources:** it can also use a wide variety of data sources, including those available in RevoScaleR

The same code will run on small and huge data, on a single core and on multiple SMP cores, on a single node and on multiple nodes.



The RevoPemaR package provides

- PemaBaseClass, which is a Reference Class
 - Pema classes inherit from (contain) this
 - This class has several methods which can be overridden by child classes. The key methods that are typically overridden are:
 - **initialize()** – initialize fields; used on construction
 - **processData()** – computes intermediate results from a chunk of data
 - **updateResults()** – updates one Pema object from another one
 - **processResults()** – converts intermediate results to final results



The RevoPemaR package also provides

- **setPemaClass()** function: a wrapper around setRefClass() to create a class generator
- **pemaCompute()** function: a function to initiate a computation
- Support for both iterative and non-iterative algorithms
- Example code:
 - PemaMean: variable mean
 - PemaWordCount, PemaPopularWords: text mining
 - PemaGradDescent, PemaLogitGD: a gradient descent base class and a logistic regression child



Overview of Pema execution

- The **initialize()** method of the master Pema object is executed
- The master Pema object is serialized and sent to each worker process
- The worker processes call **processData()** once for each chunk of data
 - The fields of the worker's Pema object are updated from the data
 - In addition, a data frame may be returned from `processData()`, and will be written to an output data source
 - When a worker has processed all of its data, it sends its reserialized Pema object back to the master (or an intermediate combiner)
- The master process loops over all of the Pema objects returned to it, calling **updateResults()** to update its Pema object
- **processResults()** is then called on the master Pema object to convert intermediate results to final results
- `hasConverged()`, whose default returns TRUE, is called, and either the results are returned to the user or another iteration is started



Typical steps in using RevoPemaR

- Write the class definition, using `setPemaClass()`, and inheriting from `RevoBaseClass` or a child class
 - Specify the fields (member variables) of the class
 - Override the `initialize()`, `processData()`, `updateResults()`, `processResults()` and any other required methods
- Test the individual methods as they are being written; data in memory (data frame, vector, matrix) can be used, or any of the `RevoScaleR` data sources
- Execute the code using the `pemaCompute()` function or the `compute()` method



An example: computing a sample mean

To create a Pema class generator function, use the `setPemaClass` function. There are four basic pieces of information that need to be specified: the class name, the super classes, the fields, and the methods

```
PemaMean <- setPemaClass(  
  Class = "PemaMean",  
  contains = "PemaBaseClass",  
  fields = list( # To be written  
    ),  
  methods = list( # To be written  
    )  
)
```



PemaMean fields

The list of field names and their types. The type “ANY” can be used to allow flexible types, but requires the author to check types

```
fields = list(  
    sum = "numeric",  
    totalObs = "numeric",  
    totalValidObs = "numeric",  
    mean = "numeric",  
    varName = "character"  
),
```



PemaMean methods: initialize

The initialize method is called when the object is constructed, and can also be called directly.

```
methods = list(  
  "initialize" = function(varName = "", ...)  
  {  
    'sum, totalValidObs, and mean are all initialized to  
    0'  
    callSuper(...) # calls the method of the parent class  
    usingMethods(.pemaMethods) # for distributed computing  
  
    # Fields are modified in a method by using the non-  
    local assignment op  
    varName <<- varName  
    sum <<- 0  
    totalObs <<- 0  
    totalValidObs <<- 0  
    mean <<- 0  
  },
```



PemaMean methods: processData

The `processData` method usually does most of the work. It takes a chunk of data and uses it to update the fields (state) of the object. It can also return a data frame of results; these will be written to an output data source.

```
"processData" = function(dataList)
{
  'Updates the sum and total observations from the current
  chunk of data.'

  if (is.null(dataList[[varName]]))
    stop( "The variable ", varName, " cannot be found
  in the data." )

  sum <-< sum + sum(as.numeric(dataList[[varName]]),
na.rm = TRUE)
  totalObs <-< totalObs + length(dataList[[varName]])
  totalValidObs <-< totalValidObs +
sum(!is.na(dataList[[varName]]))
  invisible(NULL)
},
```



PemaMean methods: updateResults

The `updateResults` method updates the fields of one Pema object from the fields of another Pema object. This is called during distributed computations to update a master object from the results of each of the worker objects. Can also be used to update yesterday's results from results obtained from today's data.

```
"updateResults" = function(pemaMeanObj)
{
  'Updates the sum and total observations from another
  PemaMean object.'
  sum <<- sum + pemaMeanObj$sum
  totalObs <<- totalObs + pemaMeanObj$totalObs
  totalValidObs <<- totalValidObs + pemaMeanObj$totalValidObs
  invisible(NULL)
},
```



PemaMean methods: processResults

The processResults method converts intermediate results in a Pema object into final results.

```
"processResults" = function()  
{  
  'Returns the sum divided by the totalValidObs.'  
  if (totalValidObs > 0)  
  {  
    mean <<- sum/totalValidObs  
  }  
  else  
  {  
    mean <<- as.numeric(NA)  
  }  
  return( mean )  
},
```



Using PemaMean

The `pemaCompute` function puts a Pema object to work.

```
# Instantiate a PemaMean object
meanPemaObj <- PemaMean()

# Call pemaCompute
mean <- pemaCompute(pemaObj = meanPemaObj,
                   data = data.frame(x = rnorm(1000)), varName = "x")

# Note that the reference class object has been updated
meanPemaObj$mean
meanPemaObj$totalValidObs

# Update meanPemaObj with new data
mean <- pemaCompute(pemaObj = meanPemaObj,
                   data = data.frame(x = rnorm(1000)), varName = "x",
                   initPema = FALSE)

meanPemaObj$totalValidObs
```



RevoScaleR Package

- C++ PEMA framework, integrated with R
- Proprietary package distributed with Revolution R Enterprise
- PEMA algorithms written in C++
- Supports distributed computing “contexts”
- Supports chunking “data sources”
- Enhances the RevoPemaR package

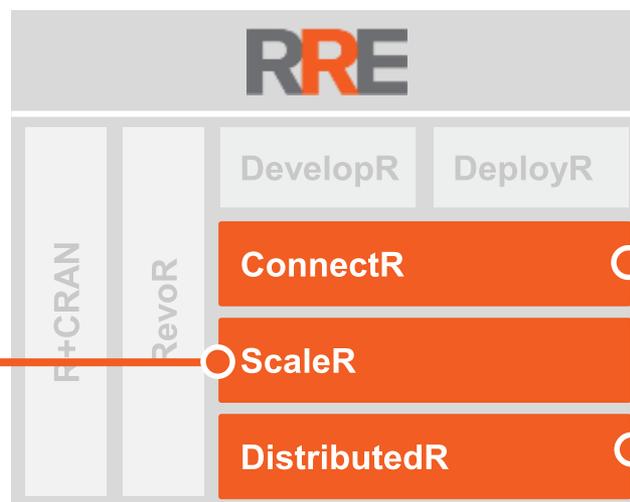


RevoScaleR:

Parallelized Algorithms, Distributed Computing, and Data Sourcing

ScaleR

- Ready-to-Use high-performance big data big analytics
- Fully-parallelized analytics
- Data prep & data distillation
- Descriptive statistics & statistical tests
- Correlation & covariance matrices
- Predictive Models – linear, logistic, GLM
- Machine learning
- Monte Carlo simulation
- Tools for distributing customized algorithms across nodes



ConnectR

- High-speed & direct connectors
- Available for:**
- High-performance XDF
 - SAS, SPSS, delimited & fixed format text data files
 - Hadoop HDFS (text & XDF)
 - Teradata Database
 - ODBC

DistributedR

- Distributed computing framework
 - Delivers portability across platforms
- Available on:**
- Teradata Database
 - Hortonworks / Cloudera / MapR
 - Windows Servers / HPC Clusters
 - IBM Platform LSF Linux Clusters
 - Red Hat Linux Servers
 - SuSE Linux Servers



Revolution R Enterprise ScaleR: *High Performance Big Data Analytics*

Data Prep, Distillation & Descriptive Analytics

R Data Step



- Data import – Delimited, Fixed, SAS, SPSS, ODBC
- **Variable creation & transformation using any R functions and packages**
- Recode variables
- Factor variables
- Missing value handling
- Sort
- Merge
- Split
- Aggregate by category (means, sums)

Descriptive Statistics



- Min / Max
- Mean
- Median (approx.)
- Quantiles (approx.)
- Standard Deviation
- Variance
- Correlation
- Covariance
- Sum of Squares (cross product matrix)
- Pairwise Cross tabs
- Risk Ratio & Odds Ratio
- Cross-Tabulation of Data
- Marginal Summaries of Cross Tabulations

Statistical Tests



- Chi Square Test
- Kendall Rank Correlation
- Fisher's Exact Test
- Student's t-Test

Sampling



- Subsample (observations & variables)
- Random Sampling



Revolution R Enterprise ScaleR (continued)

Statistical Modeling

Predictive Models



- Covariance/Correlation/Sum of Squares/Cross-product Matrix
- Multiple Linear Regression
- Logistic Regression
- Generalized Linear Models (GLM) - All exponential family distributions: binomial, Gaussian, inverse Gaussian, Poisson, Tweedie. Standard link functions including: cauchit, identity, log, logit, probit.
 - User defined distributions & link functions.
- Classification & Regression Trees and Forests
- Gradient Boosted Trees
- Residuals for all models

Data Visualization



- Histogram
- ROC Curves (actual data and predicted values)
- Lorenz Curve
- Line and Scatter Plots
- Tree Visualization

Variable Selection



- Stepwise Regression
 - Linear
 - Logistic
 - GLM

Simulation and HPC



- Monte Carlo
- **Run open source R functions and packages across cores and nodes**

Machine Learning

Cluster Analysis



- K-Means

Classification & Regression



- Decision Trees
- Decision Forests
- Gradient Boosted Trees

Deployment



- Prediction (scoring)
- PMML Export



In Conclusion

- Separate 1) processing of data from 2) the combining of intermediate results and 3) the creating of final results
- Reference Classes let you write a Pema class that inherits from RevoPemaR and can run across nodes without having to install your code on those nodes, because your code is self-contained in the Pema object
- Many options for providing open source back ends for RevoPemaR



Thank you!

- Lee Edlefsen and Sue Ranney
- John Chambers for R Reference Classes
- R-Core Team
- R Package Developers
- R Community
- Revolution R Enterprise Customers and Testers
- Colleagues at Revolution Analytics