

EARL London | 10-12 September, 2019

Large-Scale Time Series Forecasting in Apache Spark

Timothy Wong, Phuong Pham

Centrica PLC

Centrica

- We supply energy and services to over 27 million customer accounts
- Supported by around 12,000 engineers and technicians
- Our areas of focus are Energy Supply & Services, Connected Home, Distributed Energy & Power, Energy Marketing & Trading







Power Forecasting

• Account level hourly forecast for every smart meter customer in the ERCOT (Texas, USA) market













Power usage at account level (A typical example)





How the data is used?



Problems

- Forecast is required at account (meter) level
 - We have many accounts (75K)
 - Assume 1 sec per account, sequential run takes 21 hours!
- Power forecast is time-sensitive
 - Multiple runs per-day
 - Forecast must be produced on time
- Accuracy is important
 - Desire for sophisticated models
 - Sophisticated models are usually long-running!



Architecture







Lightning-fast unified analytics engine

Download Libraries - Documentation - Examples Community - Developers -

Apache Spark™ is a unified analytics engine for large-scale data processing.

- Apache Spark
 - Allow high degree of parallelism
 - Handling large datasets
 - Run native R code in Spark
- Benefits
 - Make use of sophisticated algorithms in R
 - Parallelism can save a lot of time
 - Scalable by design
- Disadvantages
 - Compute power can be expensive







Parallelism

- Apache Spark is a parallel compute framework
- Divide data into many small partitions ٠
- Increasing the number of partitions will encourage parallelism
- Too many partitions may incur excessive I/O overhead ٠







Executor 1

Task 2

Task 3

Task 4

Task 5

Task 6

AWS Elastic MapReduce (EMR)



aws Services	🗸 🔹 Resource Groups 🗸 🖈			↓ Timothy.Wong@cen	trica.com 👻 N. Virg	inia 🕶 Support 🕶					
Amazon EMR	Create cluster View details	Clone Terminate									
Clusters	Filter: All clusters Filter loaded clusters 100 clusters loaded load more C										
Security configurations	Name	ID	Status	Creation time (UTC+1) 🚽	Elapsed time	Normalized instance hours					
Events	► ► Tfs-emr-cluster-25214	j-35FPOW2BW4YS0	Running	2019-07-06 02:01 (UTC+1)	55 minutes	0					
Notebooks	tfs-emr-cluster-2521	j-2R8SH3S0ZW5MY	Running	2019-07-06 02:01 (UTC+1)	55 minutes	0					
Help	tfs-emr-cluster-2521	j-9J0JQKWT92IJ	Terminated User request	2019-07-06 02:00 (UTC+1)	55 minutes	24					
What's new	tfs-emr-cluster-25212	j-313N63VHLCK4K	Terminated User request	2019-07-06 02:00 (UTC+1)	21 minutes	576					
	tfs-emr-cluster-25210	j-2ATHAK01U7XRE	Terminated User request	2019-07-06 01:01 (UTC+1)	41 minutes	352					
	tfs-emr-cluster-25211	j-HIDQ2ILQDR5T	Terminated User request	2019-07-06 01:01 (UTC+1)	33 minutes	352					
	mRapid_EMR_122	j-IRG9PMSZYVTQ	Terminated User request	2019-07-06 01:01 (UTC+1)	1 hour, 9 minutes	24					
	tfs-emr-cluster-25209	j-10J09AZJV56Q5	Terminated User request	2019-07-06 01:00 (UTC+1)	1 hour, 40 minutes	704					
	tfs-emr-cluster-25208	j-1YWED20503B1E	Terminated User request	2019-07-06 00:00 (UTC+1)	26 minutes	24					
	tfs-emr-cluster-25207	i-315D3VVOSIGFI	Terminated	2019-07-05 23:05 (UTC+1)	38 minutes	160					



Setting up EMR

- EC2 charge + EMR charge
- Billed per second
- Terminate once job is completed
- Pay for what you use

Pricing for Amazon EMR and Amazon EC2 (On-Demand)

Region: US East (N. Virginia) +

	Amazon EC2 Price	Amazon EMR Price
General Purpose - Current Generation		
m5.xlarge	\$0.192 per Hour	\$0.048 per Hour
m5.2xlarge	\$0.384 per Hour	\$0.096 per Hour
m5.4xlarge	\$0.768 per Hour	\$0.192 per Hour
m5.12xlarge	\$2.304 per Hour	\$0.27 per Hour
m5.24xlarge	\$4.608 per Hour	\$0.27 per Hour
m5a.xlarge	\$0.172 per Hour	\$0.043 per Hour
mEa Ovlargo	¢0 Z44 por Hour	¢0.0% por Hour

EARL²

CONFERENCE

Cloud Architecture







Connecting to Spark (EMR)

- 37 library(sparklyr)
- 38 conf <- spark_config()</pre>
- 39 conf["sparklyr.log.console"] <- TRUE</pre>
- 40 conf["spark.network.timeout"] <- "3600s"</pre>
- 41 conf["spark.sql.crossJoin.enabled"] <- "true"</pre>
- 42 conf["spark.sql.parquet.compression.codec"] <- "snappy"</pre>
- 43 conf["spark.r.command"] <- "/opt/R/3.5.0/lib64/R/bin/Rscript"</pre>
- 44 conf["spark.sql.shuffle.partitions"] <- "8000"</pre>
- 45 conf["spark.default.parallelism"] <- "8000"</pre>
- 46 conf["spark.scheduler.listenerbus.eventqueue.capacity"] <- "100000"</pre>
- 47 conf["spark.dynamicAllocation.enabled"] <- "true"</pre>
- 48 conf["spark.shuffle.service.enabled"] <- "true"</pre>
- 49 conf["spark.driver.maxResultSize"] <- "32G"</pre>
- 50 conf["spark.serializer"] <- "org.apache.spark.serializer.KryoSerializer"</pre>
- 51 conf["spark.kryoserializer.buffer.max"] <- "2047m"</pre>
- 52 conf["sparklyr.apply.serializer"] <- "org.apache.spark.serializer.KryoSerializer"
 53</pre>

EARL

CONFERENCE 🔁

- 54 sc <- spark_connect(master = "yarn",
- 55 config = conf)

Resource Manager



NEW, NEW_SAVING, SUBMITTED, ACCEPTED, RUNNING Applications

* Cluster	Cluster Metrics									
About Nodes	Apps Submitted Apps Pe 5 0	ending Apps Running 1	Apps Completed 4	Containers Running 31	Memory Used 1.03 TB 1	Memory Total .17 TB 1-	Memory Reserved 40.88 GB	VCores Used 31	VCores Total VCore 160 4	s Reserved
Node Labels	Cluster Nodes Metrics									
NEW	Active Nodes	Decommissioning Nodes	De	ecommissioned Nodes	Lost Nodes	Unhe	ealthy Nodes	Rebooted Nodes	s Shutdown	Nodes
NEW SAVING SUBMITTED	<u>10</u> <u>0</u> Scheduler Metrics		Q		<u>0</u>	<u>0</u>	<u>C</u>	1	Q	
ACCEPTED RUNNING	Scheduler Type	Scheduling Resou	ce Type	Minimum Allocation		Maximum Allo	cation	Maxim	um Cluster Application Priority	
FINISHED	Capacity Scheduler	[MEMORY]	<me< td=""><td>emory:32, vCores:1></td><td><memory:< td=""><td>122880, vCores:16></td><td></td><td>0</td><td>•</td><td></td></memory:<></td></me<>	emory:32, vCores:1>	<memory:< td=""><td>122880, vCores:16></td><td></td><td>0</td><td>•</td><td></td></memory:<>	122880, vCores:16>		0	•	
KILLED	Dump scheduler logs 1 min 🔻									
Scheduler	Application Queues									
Tools	Legend: Capacity	Used Over	apacity) Max	Capacity Users F	Requesting Resource	es				
	- Partition: <default_partitio< td=""><td>N> <memory:0, vcores:0=""></memory:0,></td><td></td><td></td><td></td><td></td><td></td><td></td><td>0.0% used</td><td></td></default_partitio<>	N> <memory:0, vcores:0=""></memory:0,>							0.0% used	
	+ Queue: root								0.0% used	
	Partition: CORE <memory:122< td=""><td>8800, vCores:160></td><td></td><td></td><td></td><td></td><td></td><td></td><td>99.9% used</td><td></td></memory:122<>	8800, vCores:160>							99.9% used	
	Show 20 • entries							\frown	Search:	
	ID \$	User Name A	pplication Queue Appl Type ≎ ≎ Pric	lication StartTime FinishTi	me State o Finals	Status Containers	Allocated Allocated CPU Memory VCores © MB	% of % of Queue Cluster I	Progress © Tracking UI	Blacklisted Nodes ©
	application_1562358292204_0005	radmin step_02_modelling SF	ARK efault 0	Sat Jul 6 02:40:00 +0100 2019	RUNNING UNDE	FINED 31	31 108281	88.1 88.1	ApplicationMaste	<u>er</u> 0
	Showing 1 to 1 of 1 entries							4	First Previous	Next Last
Sc	oark job is rui	nning in YA	RN			<u>Resour</u>	ce is allo	ocated t	<u>o this job</u>	



Logged in as: dr.who

Spark UI

A	
Sood	
SOUIN	240

Jobs Stages

Storage Environment Executors SQL

step_02_modelling application UI

Spark Jobs (?)

User: radmin Total Uptime: 16 min Scheduling Mode: FIFO Active Jobs: 1 Completed Jobs: 8

Event Timeline

- Active Jobs (1)

Job Id 🔹	Description		Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
8	csv at NativeMethodAccessorImpl.java:0 csv at NativeMethodAccessorImpl.java:0	(kill)	2019/07/06 01:48:07	7.9 min	1/2	1005/1205 (1 running)

- Completed Jobs (8)

Job Id	Description	Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
7	collect at utils.scala:204 collect at utils.scala:204	2019/07/06 01:47:42	21 s	2/2	1006/1006
6	collect at utils.scala:204 collect at utils.scala:204	2019/07/06 01:47:16	63 ms	1/1	1/1
5	collect at utils.scala:204 collect at utils.scala:204	2019/07/06 01:47:11	0.5 s	2/2	1006/1006
4	sql at NativeMethodAccessorImpl.java:0 sql at NativeMethodAccessorImpl.java:0	2019/07/06 01:43:16	3.9 min	2/2	1006/1006
3	csv at NativeMethodAccessorImpl.java:0 csv at NativeMethodAccessorImpl.java:0	2019/07/06 01:40:24	2.8 min	1/1	1005/1005



Running R in Spark

```
30
    my_tbl_lcl <- data.frame(x = rnorm(100),
31
                              y = rnorm(100),
32
                              group = c(1,2)
33
    my_tbl <- sc %>% copy_to(my_tbl_lcl)
34
35
    my_results <- my_tbl %>% spark_apply(function(df){
36 -
37
      my model <- lm(y \sim x, df)
38
      return(my_model$residuals)
   }, group_by = "group")
39
```



Model Training / Deployment

- Training a new model is time consuming
- Batch job to retrain all models regularly
 - Retrain all accounts
 - Takes ~20 hours on AWS EMR
 - 1,200 vCPU & 9.6 TB memory
- Model objects are persisted
 - Saved to AWS S3 as compressed .rda file
 - Can be loaded back later for deployment
 - I/O to S3 is operated on Spark worker nodes







Loading R object from S3

- 1 library(reticulate)
- 2 boto3 <- import("boto3")</pre>
- 3 s3 <- boto3\$resource("s3")
- 4 my_file <- s3\$Object("my-bucket-name","my_folder/my_file.RData")</pre>
- 5 my_tmp <- tempfile()</pre>
- 6 my_file\$download_file(my_tmp)
- 7 my_object <- readRDS(my_tmp)</pre>
- 8 unlink(my_tmp)



Effects of Partitioning

- The execution time of the tasks in the longest stage (*spark_apply*) is measured
 - Timed using r5.4xlarge with 10 worker nodes
- This stage governs the total execution time of the pipeline
- There is additional overhead time for data preparation (trivial) and I/O to S3

Boxplot showing the effects of partitioning on task execution time



Duration (Minutes) Partitions Models 5% 25% Median 75% 95% 500 396 432 462 490 519 64 84 96 105 126 2000 LM only 4000 24 36 42 54 66 8000 9 17 28 22 38 8000 18 84 LM + SARIMA + SARIMAX 40 59 119 8000 LM + RPART + RF + GBM 16 26 41 53 66

EARL SCONFERENCE

Modelling



Multi-model approach

- Power usage at account level have different behaviours
 - Some are sensitive to weather (e.g. heating / cooling systems)
 - Some have weekly patterns (e.g. office building)
 - Some have regular behaviour (e.g. warehouse)
 - Some have predictable spikes (e.g. cooking at lunch time)
 - Some have unpredictable spikes (?)
 - •
- We need different modelling techniques
 - Run several models for *all* accounts



Algorithms

Acronym	Algorithm	Description
GLM	Generalised Linear Model	Linear regression
SARIMA	Seasonal Autoregressive Integrative Moving Average	Time series model which can deal with seasonality
SARIMAX	Seasonal Autoregressive Integrative Moving Average with Exogeneous Variables	Same as SARIMA but can handle covariate variables (e.g. temperature)
MSARIMA	Multi-Seasonal Autoregressive Integrative Moving Average	Same as SARIMA but handle multiple seasonalities (e.g. daily, weekly, annual)
MSARIMAX	Multi-Seasonal Autoregressive Integrative Moving Average with Exogeneous Variable	Same as MSARIMA but can handle covariate variables.
RPART	Recursive Partitioning	Decision tree model grows deep and pruned afterwards.
RF	Random Forest	Ensemble of large decision trees with randomly selected variable at each split.
GBM	Gradient Boosting Machine	Shallow decision trees are introduced additively, forming a strong learner.
ELASTICNET	Linear Regression with ElasticNet Regularisation	Regularised linear regression



Dynamic Weighting





Results



Results - Backcast

- Models were trained using 2 years worth of data up to the end of March 2019
- Backcast is produced using actualised weather data in April 2019 (full month)
- Chart below shows the backcast results (Solid = backcast; Dotted = actual)



Results - Forecast

• 10 days ahead forecast, aggregated at zonal level





Accuracy

- The MAPE of each model were calculated for the 1 month backcast period
- Tree-based methods (GBM, RF, RPART) outperform regression methods
 - Likely due to the discovery of interaction effects
 - Deals with non-linearity very well
- Time series methods (SARIMA, SARIMAX, MSARIMA, MSARIMAX) have worst accuracy
 - Account level usage may have many change points and intermittent spikes, rendering time series methods less useful.

Model	HOUSTON	NORTH	SOUTH	WEST
GBM	5.67	5.83	4.86	3.02
RPART	6.24	5.96	5.10	3.56
RF	6.90	7.20	5.93	3.84
ELASTICNET	9.68	6.54	7.19	4.18
GLM	7.11	7.20	7.16	4.37
MSARIMA	16.97		18.69	10.68
MSARIMAX	17.69		19.59	11.24
SARIMA	35.16		36.97	14.24
SARIMAX	49.35		97.26	38.27



Summary



Technologies Used





Potential Improvements





Q & A

Timothy Wong CSCI CStat CEng MBCS Principal Data Scientist, Centrica plc



timothy.wong@centrica.com



linkedin.com/in/timothywong731



@timothywong731



timothywong731.github.io

Phuong Pham

Data Scientist, Centrica plc



phuong.pham@centrica.com



linkedin.com/in/phuong-pham-87498251

