

# Data-Intensive Distributed Computing

CS 431/631 451/651 (Fall 2019)

Part 2: From MapReduce to Spark (1/2)

September 19, 2019

Ali Abedi


These slides are available at <http://roegiest.com/bigdata-2019w/>

This work is licensed under a Creative Commons Attribution-Noncommercial-Share Alike 3.0 United States  
See <http://creativecommons.org/licenses/by-nc-sa/3.0/us/> for details



An aerial photograph of a large industrial datacenter facility during sunset. The sun is low on the horizon, casting a warm orange glow over the scene. The facility consists of several large, white, rectangular buildings with flat roofs, arranged in a grid-like pattern. In the foreground, there is a large parking lot filled with many white semi-trailers. The surrounding area is a mix of green fields and brown, tilled soil. The text "The datacenter is the computer! What's the instruction set?" is overlaid in white, sans-serif font across the middle of the image.

The datacenter *is* the computer!  
What's the instruction set?



So you like programming in assembly?

# What's the solution?

Design a higher-level language

Write a compiler

Hadoop is great, but it's really waaaaay too low level!



What we really need is SQL!

Answer:



What we really need is a scripting language!


Answer:



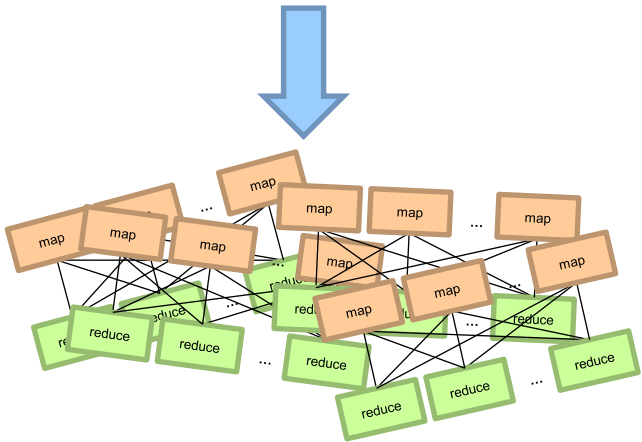


SQL

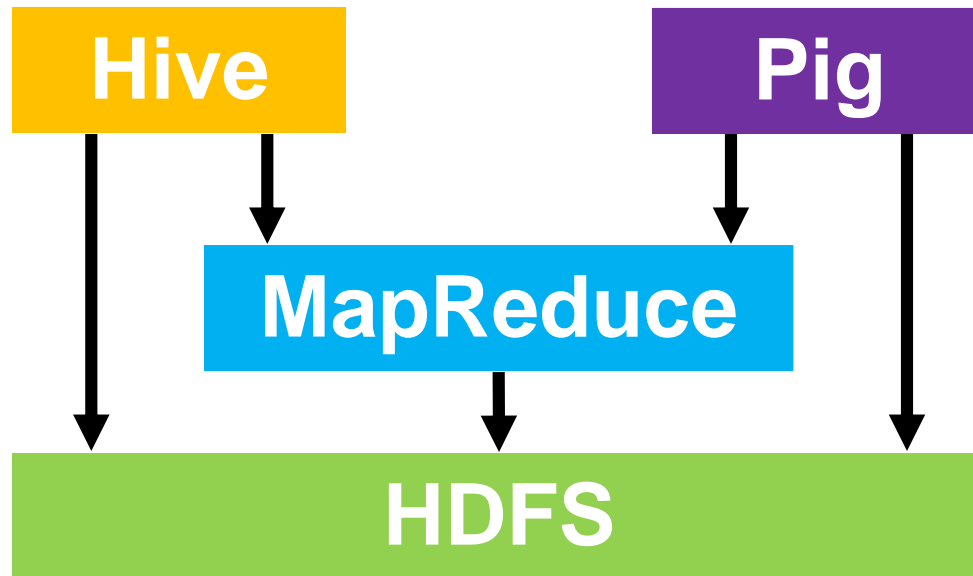
Aside: Why not just use a database?



Pig Scripts



Both open-source projects today!





Pig!



# Pig: Example

Task: Find the top 10 most visited pages in each category

Visits

User	Url	Time
Amy	cnn.com	8:00
Amy	bbc.com	10:00
Amy	flickr.com	10:05
Fred	cnn.com	12:00



URL Info

Url	Category	PageRank
cnn.com	News	0.9
bbc.com	News	0.8
flickr.com	Photos	0.7
espn.com	Sports	0.9

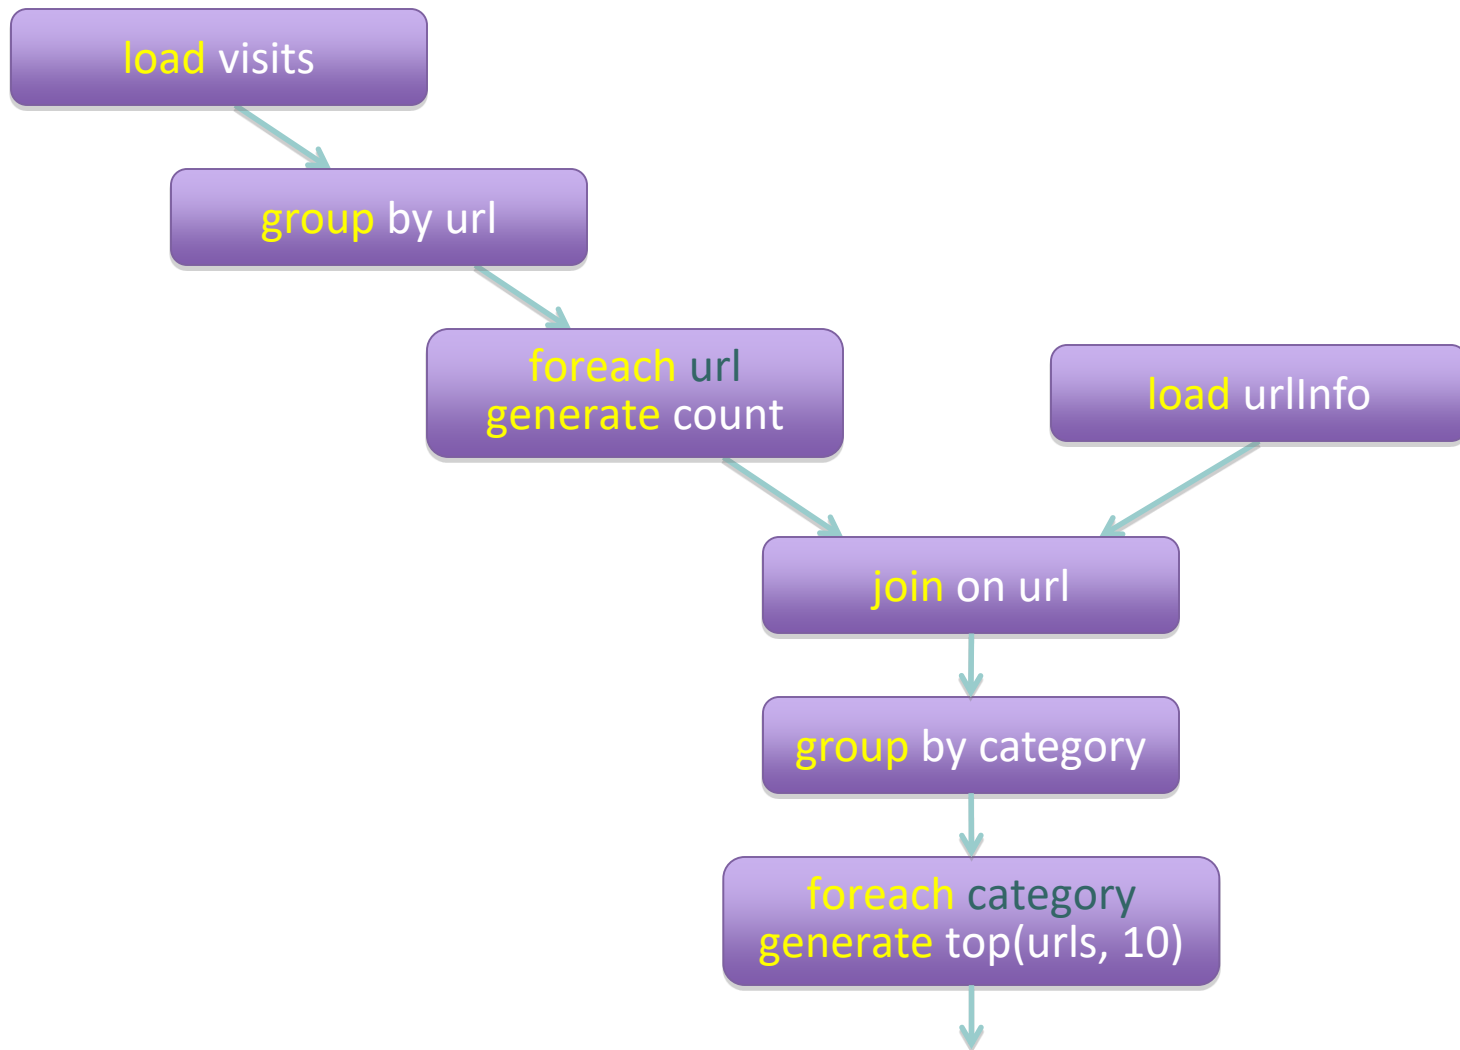


# Pig: Example Script

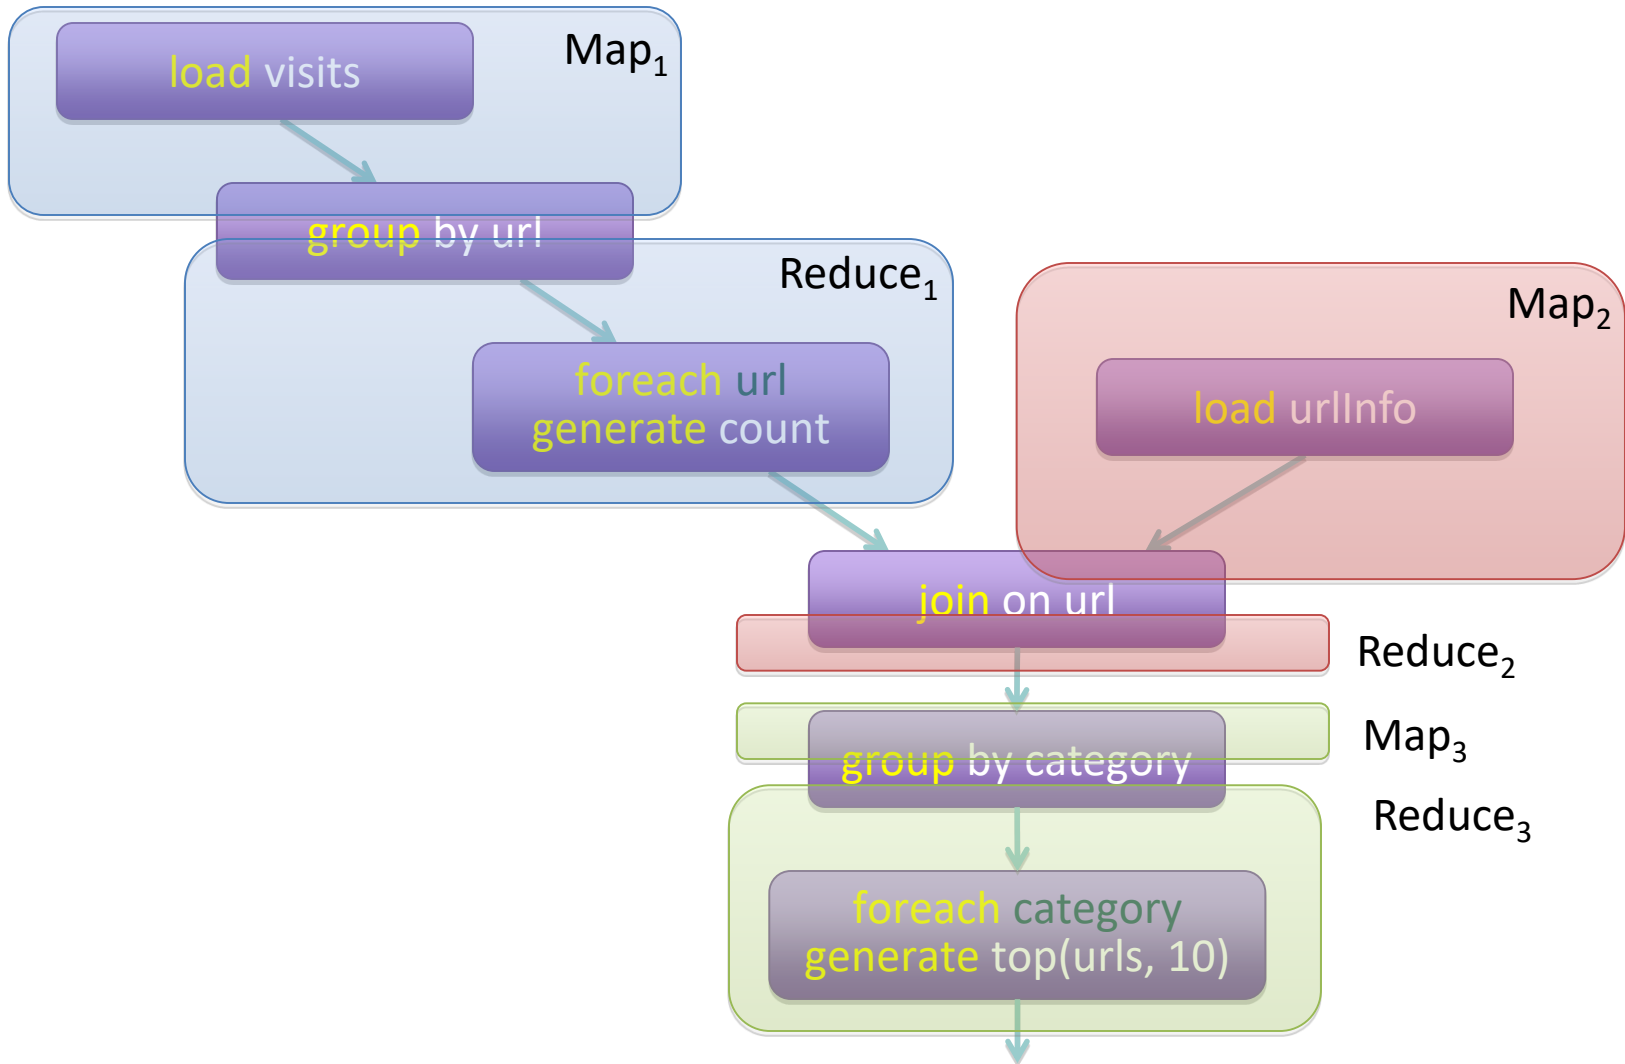
```
visits = load '/data/visits' as (user, url, time);
gVisits = group visits by url;
visitCounts = foreach gVisits generate url, count(visits);
urlInfo = load '/data/urlInfo' as (url, category, pRank);
visitCounts = join visitCounts by url, urlInfo by url;
gCategories = group visitCounts by category;
topUrls = foreach gCategories generate top(visitCounts,10);

store topUrls into '/data/topUrls';
```

# Pig Query Plan



# Pig: MapReduce Execution



visits = load '/data/visits' as (user, url, time);  
 gVisits = group visits by url;  
 visitCounts = foreach gVisits generate url, count(visits);  
 urlInfo = load '/data/urlInfo' as (url, category, pRank);  
 visitCounts = join visitCounts by url, urlInfo by url;  
 gCategories = group visitCounts by category;  
 topUrls = foreach gCategories generate top(visitCounts,10);

This?

store topUrls into '/data/topUrls';

```

import java.io.IOException;
import java.util.ArrayList;
import java.util.Iterator;
import java.util.List;

import org.apache.hadoop.fs.Path;
import org.apache.hadoop.io.LongWritable;
import org.apache.hadoop.io.Text;
import org.apache.hadoop.io.Writable;
import org.apache.hadoop.io.WritableComparable;
import org.apache.hadoop.mapred.FileInputFormat;
import org.apache.hadoop.mapred.FileOutputFormat;
import org.apache.hadoop.mapred.JobConf;
import org.apache.hadoop.mapred.KeyValueTextInputFormat;
import org.apache.hadoop.mapred.Mapper;
import org.apache.hadoop.mapred.MapReduceBase;
import org.apache.hadoop.mapred.OutputCollector;
import org.apache.hadoop.mapred.RecordReader;
import org.apache.hadoop.mapred.Reducer;
import org.apache.hadoop.mapred.Reporter;
import org.apache.hadoop.mapred.SequenceFileInputFormat;
import org.apache.hadoop.mapred.SequenceFileOutputFormat;
import org.apache.hadoop.mapred.TextInputFormat;
import org.apache.hadoop.mapred.JobControl.JobC
import org.apache.hadoop.mapred.lib.IdentityMapper;

public class MRExample {
    public static class LoadPages extends MapReduceBase
    implements Mapper<LongWritable, Text, Text, Text> {

        public void map(LongWritable k, Text val,
            OutputCollector<Text, Text> oc,
            Reporter reporter) throws IOException {
            // Pull the key out
            String line = val.toString();
            int firstComma = line.indexOf(',');
            String key = line.substring(0, firstComma);
            String value = line.substring(firstComma + 1);
            Text outKey = new Text(key);
            // Append an index to the value so we know which file
            // it came from.
            Text outVal = new Text("1" + value);
            oc.collect(outKey, outVal);
        }

        public static class LoadAndFilterUsers extends MapReduceBase
        implements Mapper<LongWritable, Text, Text, Text> {

            public void map(LongWritable k, Text val,
                OutputCollector<Text, Text> oc,
                Reporter reporter) throws IOException {
                // Pull the key out
                String line = val.toString();
                int firstComma = line.indexOf(',');
                String value = line.substring(
                    firstComma + 1);
                int age = Integer.parseInt(value);
                if (age < 18 || age > 25) return;
                String key = line.substring(0, firstComma);
                Text outKey = new Text(key);
                // Prepend an index to the value so we know which file
                // it came from.
                Text outVal = new Text("2" + value);
                oc.collect(outKey, outVal);
            }

            public static class Join extends MapReduceBase
            implements Reducer<Text, Text, Text, Text> {

                public void reduce(Text key,
                    Iterator<Text> iter,
                    OutputCollector<Text, Text> oc,
                    Reporter reporter) throws IOException {
                    // For each value, figure out which file it's from and
                    store it
                    reporter.setStatus("OK");
                }

                // Do the cross product and collect the values
                for (String s1 : first) {
                    for (String s2 : second) {
                        String outVal = key + "/" + s1 + "/" + s2;
                        oc.collect(null, new Text(outVal));
                        reporter.setStatus("OK");
                    }
                }

                public static class LoadJoined extends MapReduceBase
                implements Mapper<Text, Text, Text, LongWritable> {

                    public void map(
                        Text k,
                        Text val,
                        OutputCollector<Text, LongWritable> oc,
                        Reporter reporter) throws IOException {
                        // Find the url
                        String line = val.toString();
                        int firstComma = line.indexOf(',');
                        String key = line.substring(firstComma, secondComma);
                        // drop the rest of the record, I don't need it anymore,
                        // just pass a 1 for the combiner/reducer to sum instead.
                        Text outKey = new Text(key);
                        oc.collect(outKey, new LongWritable(1));
                    }

                    public static class ReduceURLs extends MapReduceBase
                    implements Reducer<Text, LongWritable, WritableComparable,
                    Writable> {

                        public void reduce(
                            Text k,
                            Iterator<LongWritable> iter,
                            OutputCollector<WritableComparable, Writable> oc,
                            Reporter reporter) throws IOException {
                                // Add up all the values we see
                                long sum = 0;
                                while (iter.hasNext()) {
                                    sum += iter.next().get();
                                    reporter.setStatus("OK");
                                }
                                oc.collect(key, new LongWritable(sum));
                            }

                            public static class LoadClicks extends MapReduceBase
                            implements Mapper<WritableComparable, Writable, LongWritable,
                            Text> {

                                public void map(
                                    WritableComparable key,
                                    Writable val,
                                    OutputCollector<LongWritable, Text> oc,
                                    Reporter reporter) throws IOException {
                                        oc.collect((LongWritable)val, (Text)key);
                                    }

                                    public static class LimitClicks extends MapReduceBase
                                    implements Reducer<LongWritable, Text, LongWritable, Text> {

                                        int count = 0;
                                        public void reduce(
                                            LongWritable key,
                                            Iterator<Text> iter,
                                            OutputCollector<LongWritable, Text> oc,
                                            Reporter reporter) throws IOException {
                                                // Only output the first 100 records
                                                reporter.setStatus("OK");
                                            }

                                            JobConf ifu = new JobConf(MRExample.class);
                                            ifu.setJobName("Load and Filter Users");
                                            ifu.setInputFormat(TextInputFormat.class);
                                            ifu.setOutputKeyClass(Text.class);
                                            ifu.setOutputValueClass(Text.class);
                                            ifu.setMapperClass(LoadAndFilterUsers.class);
                                            FileInputFormat.addInputPath(ifu, new
                                            Path("/user/gates/users"));
                                            FileOutputFormat.setOutputPath(ifu, new
                                            Path("/user/gates/tap/filtered_users"));
                                            ifu.setNumReduceTasks(0);
                                            Job loadUsers = new Job(ifu);

                                            JobConf join = new JobConf(
                                                MRExample.class);
                                            join.setJobName("Join Users and Pages");
                                            join.setInputFormat(KeyValueTextInputFormat.class);
                                            join.setOutputKeyClass(Text.class);
                                            join.setOutputValueClass(Text.class);
                                            join.setMapperClass(IdentityMapper
                                                .class);
                                            join.setReducerClass(Join.class);
                                            FileInputFormat.addInputPath(join, new
                                            Path("/user/gates/tap/indexed_pages"));
                                            FileInputFormat.addInputPath(join, new
                                            Path("/user/gates/tap/filtered_users"));
                                            FileOutputFormat.setOutputPath(join, new
                                            Path("/user/gates/tap/joined"));
                                            join.setNumReduceTasks(50);
                                            Job joinJob = new Job(join);
                                            joinJob.addDependingJob(loadPages);
                                            joinJob.addDependingJob(loadUsers);

                                            JobConf group = new JobConf(MRE
                                                xample.class);
                                            group.setJobName("Group URLs");
                                            group.setInputFormat(KeyValueTextInputFormat.class);
                                            group.setOutputKeyClass(Text.class);
                                            group.setOutputValueClass(LongWritable.class);
                                            group.setOutputFormat(SequenceFi
                                                leOutputFormat.class);
                                            group.setMapperClass(LoadJoined.class);
                                            group.setCombinerClass(ReduceURLs.class);
                                            group.setReducerClass(ReduceURLs.class);
                                            FileInputFormat.addInputPath(group, new
                                            Path("/user/gates/tap/joined"));
                                            FileOutputFormat.setOutputPath(group, new
                                            Path("/user/gates/tap/grouped"));
                                            group.setNumReduceTasks(50);
                                            Job groupJob = new Job(group);
                                            groupJob.addDependingJob(joinJob);

                                            JobConf top100 = new JobConf(MRExample.class);
                                            top100.setJobName("Top 100 sites");
                                            top100.setInputFormat(SequenceFileInputFormat.class);
                                            top100.setOutputKeyClass(LongWritable.class);
                                            top100.setOutputValueClass(Text.class);
                                            top100.setOutputFormat(SequenceFileOutputF
                                                ormat.class);
                                            top100.setMapperClass(LoadClicks.class);
                                            top100.setCombinerClass(LimitClicks.class);
                                            top100.setReducerClass(LimitClicks.class);
                                            FileInputFormat.addInputPath(top100, new
                                            Path("/user/gates/tap/grouped"));
                                            FileOutputFormat.setOutputPath(top100, new
                                            Path("/user/gates/top100sitesforusers18to25"));
                                            top100.setNumReduceTasks(1);
                                            Job limit = new Job(top100);
                                            limit.addDependingJob(groupJob);
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
    }
}

```

Or this?

# But isn't Pig slower?

Sure, but c can be slower than assembly too...



# Pig: Basics

Sequence of statements manipulating relations

Data model

atoms

tuples

bags

maps

# Pig: Common Operations

LOAD: load data (from HDFS)

FOREACH ... GENERATE: per tuple processing

FILTER: discard unwanted tuples *“map”*

*“reduce”* GROUP/COGROUP: group tuples

JOIN: relational join

STORE: store data (to HDFS)



# Pig: GROUPing

```
A = LOAD 'myfile.txt' AS (f1: int, f2: int, f3: int);
```

```
(1, 2, 3)
```

```
(4, 2, 1)
```

```
(8, 3, 4)
```

```
(4, 3, 3)
```

```
(7, 2, 5)
```

```
(8, 4, 3)
```

```
X = GROUP A BY f1;
```

```
(1, {(1, 2, 3)})
```

```
(4, {(4, 2, 1), (4, 3, 3)})
```

```
(7, {(7, 2, 5)})
```

```
(8, {(8, 3, 4), (8, 4, 3)})
```

# Pig: COGROUing

A:

(1, 2, 3)

(4, 2, 1)

(8, 3, 4)

(4, 3, 3)

(7, 2, 5)

(8, 4, 3)

B:

(2, 4)

(8, 9)

(1, 3)

(2, 7)

(2, 9)

(4, 6)

(4, 9)

X = COGROU A BY \$0, B BY \$0;

(1, {(1, 2, 3)}, {(1, 3)})

(2, {}, {(2, 4), (2, 7), (2, 9)})

(4, {(4, 2, 1), (4, 3, 3)}, {(4, 6), (4, 9)})

(7, {(7, 2, 5)}, {})

(8, {(8, 3, 4), (8, 4, 3)}, {(8, 9)})

# Pig: JOINing

A:

(1, 2, 3)

(4, 2, 1)

(8, 3, 4)

(4, 3, 3)

(7, 2, 5)

(8, 4, 3)

B:

(2, 4)

(8, 9)

(1, 3)

(2, 7)

(2, 9)

(4, 6)

(4, 9)

X = JOIN A BY \$0, B BY \$0;

(1,2,3,1,3)

(4,2,1,4,6)

(4,3,3,4,6)

(4,2,1,4,9)

(4,3,3,4,9)

(8,3,4,8,9)

(8,4,3,8,9)

# Pig UDFs

User-defined functions:

Java

Python

JavaScript

Ruby

...

UDFs make Pig arbitrarily extensible

Express “core” computations in UDFs

Take advantage of Pig as glue code for scale-out plumbing


An aerial photograph of a large industrial datacenter facility during sunset. The sun is low on the horizon, casting a warm orange glow over the scene. The facility consists of several large, white, rectangular buildings with flat roofs, arranged in a grid-like pattern. A prominent feature is a large, open-air area with rows of white, cylindrical cooling towers or condensers. The surrounding landscape is a mix of green fields and brown, tilled soil. In the background, there are more buildings and a road. The overall scene is a mix of industrial and agricultural elements.

The datacenter *is* the computer!

What's the instruction set?



Okay, let's fix this!

# MapReduce Workflows





What's wrong?


# Want MM?



# Want MRR?





An aerial photograph of a large industrial datacenter facility during sunset. The sun is a bright orange orb in the upper left, casting a warm glow over the scene. The facility consists of several large, white, rectangular buildings with flat roofs, arranged in a grid-like pattern. A prominent building in the lower right is filled with rows of server racks. A parking lot in the foreground is filled with many white semi-trailers. The surrounding landscape is a mix of green fields and brown agricultural land, with some smaller buildings and roads visible in the distance.

The datacenter *is* the computer!  
Let's enrich the instruction set!

# Spark

Answer to “What’s beyond MapReduce?”


## Brief history:

Developed at UC Berkeley AMPLab in 2009

Open-sourced in 2010

Became top-level Apache project in February 2014

# Spark vs. Hadoop



Google Trends