

Big Data Analytics

C. Distributed Computing Environments / C.1 Map Reduce

Lars Schmidt-Thieme

Information Systems and Machine Learning Lab (ISMLL)
Institute for Computer Science
University of Hildesheim, Germany

Syllabus

- | | | |
|------------|------|---|
| Tue. 10.4. | (1) | 0. Introduction |
| | | A. Parallel Computing |
| Tue. 17.4. | (2) | A.1 Threads |
| Tue. 24.4. | (3) | A.2 Message Passing Interface (MPI) |
| Tue. 1.5. | — | — <i>Labour Day</i> — |
| Tue. 8.5. | (4) | A.3 Graphical Processing Units (GPUs) |
| Tue. 15.5. | (5) | (ctd.) |
| Tue. 22.5. | — | — <i>Pentecoste Break</i> — |
| | | B. Distributed Storage |
| Tue. 29.5. | (6) | B.1 Distributed File Systems |
| Tue. 5.6. | (7) | B.2 Partitioning of Relational Databases |
| Tue. 12.6. | (8) | B.3 NoSQL Databases |
| | | C. Distributed Computing Environments |
| Tue. 19.6. | (9) | C.1 Map-Reduce |
| Tue. 26.6. | (10) | C.2 Resilient Distributed Datasets (Spark) |
| Tue. 3.7. | (11) | C.3 Computational Graphs (TensorFlow) |
| | | D. Distributed Machine Learning Algorithms |
| Tue. 10.7. | (12) | D.1 Distributed Stochastic Gradient Descent |

Outline

1. Introduction
2. Parallel Computing Speedup
3. Example: Counting Words
4. Map-Reduce
5. Map-Reduce Tutorial

Outline

1. Introduction
2. Parallel Computing Speedup
3. Example: Counting Words
4. Map-Reduce
5. Map-Reduce Tutorial

Technology Stack

Part D

Distributed Machine Learning Algorithms

Part C

Distributed Execution Environments

Part B

Distributed Storage


Part A

Parallel/Distributed Computing


Why do we need a Computational Model?

- ▶ Our data is nicely stored in a distributed infrastructure
- ▶ We have a number of computers at our disposal
- ▶ We want our analytics software to take advantage of all this computing power
- ▶ When programming we want to focus on understanding our data and not our infrastructure

Shared Memory Infrastructure



Distributed Infrastructure




Outline

1. Introduction
2. Parallel Computing Speedup
3. Example: Counting Words
4. Map-Reduce
5. Map-Reduce Tutorial

Parallel Computing / Speedup

- ▶ We have p processors available to execute a task T
- ▶ Ideally: the more processors the faster a task is executed
- ▶ Reality: synchronisation and communication costs
- ▶ Speedup $s(T, p)$ of a task T by using p processors:
 - ▶ Be $t(T, p)$ the time needed to execute T using p processors
 - ▶ **Speedup** is given by:

$$s(T, p) = \frac{t(T, 1)}{t(T, p)}$$



Parallel Computing / Efficiency

- ▶ We have p processors available to execute a task T
- ▶ **Efficiency** $e(T, p)$ of a task T by using p processors:

$$e(T, p) = \frac{t(T, 1)}{p \cdot t(T, p)}$$

Considerations

- ▶ It is not worth using a lot of processors for solving small problems
- ▶ Algorithms should increase efficiency with problem size

Outline

1. Introduction
2. Parallel Computing Speedup
3. Example: Counting Words
4. Map-Reduce
5. Map-Reduce Tutorial

Word Count Example

- ▶ Given a corpus of text documents

$$D := \{d_1, \dots, d_n\}$$

each containing a sequence of words:

$$(w_1, \dots, w_m)$$

from a set W of possible words.

- ▶ the task is to generate word counts for each word in the corpus

Paradigms — Shared Memory

- ▶ All the processors have access to all counters
- ▶ Counters can be overwritten
- ▶ Processors need to lock counters before using them

Shared vector for word counts: $c \in \mathbb{N}^{|W|}$
 $c \leftarrow \{0\}^{|W|}$

Each processor:

1. access a document $d \in D$
2. for each word w in document d :
 - 2.1 lock(c_w)
 - 2.2 $c_w \leftarrow c_w + 1$
 - 2.3 unlock(c_w)

Paradigms — Shared Memory

- ▶ inefficient due to waiting times for the locks
 - ▶ the more processors, the less efficient
- ▶ in a distributed scenario even worse due to communication overhead for acquiring/releasing the locks

Paradigms — Message passing

- ▶ Each processor sees only one part of the data

$$\pi(D, p) := \{d_{p \frac{n}{p}}, \dots, d_{(p+1) \frac{n}{p} - 1}\}$$

- ▶ Each processor works on its partition
- ▶ Results are exchanged between processors (message passing)

For each processor p :

1. For each $d \in \pi(D, p)$
 - 1.1 process(d)
2. Communicate results

Word Count — Message passing

We need to define two types of processes:

1. **slave**
 - ▶ counts the words on a subset of documents and informs the master
2. **master**
 - ▶ gathers counts from the slaves and sums them up

Word Count — Message passing

Slave:

Local memory:

subset of documents: $\pi(D, p) := \{d_{p \frac{n}{p}}, \dots, d_{(p+1) \frac{n}{p} - 1}\}$

address of the master: `addr_master`

local word counts: $c \in \mathbb{R}^{|W|}$

1. $c \leftarrow \{0\}^{|W|}$
2. for each document $d \in \pi(D, p)$
for each word w in document d :
 $c_w \leftarrow c_w + 1$
3. **Send message** `send(addr_master, c)`

Word Count — Message passing

Master:

Local memory:

1. **Global word counts:** $c^{\text{global}} \in \mathbb{R}^{|W|}$
2. **List of slaves:** S

$$c^{\text{global}} \leftarrow \{0\}^{|W|}$$

$$s \leftarrow \{0\}^{|S|}$$

For each received message (p, c^p)

1. $c^{\text{global}} \leftarrow c^{\text{global}} + c^p$
2. $s_p \leftarrow 1$
3. if $\|s\|_1 = |S|$ return c^{global}

Paradigms — Message passing

- ▶ We need to manually assign master and slave roles for each processor
- ▶ Partition of the data needs to be done manually
- ▶ Implementations like OpenMPI only provide services to exchange messages

Outline

1. Introduction
2. Parallel Computing Speedup
3. Example: Counting Words
- 4. Map-Reduce**
5. Map-Reduce Tutorial

Map-Reduce

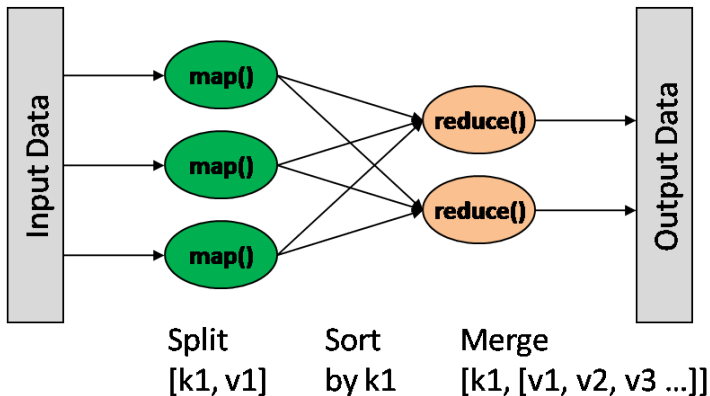
- ▶ distributed computing environment
 - ▶ introduced 2004 by Google
 - ▶ open source reference implementation: **Hadoop** (since 2006)
 - ▶ meanwhile supported by many distributed programming environments
 - ▶ e.g., in document databases such as **MongoDB**
- ▶ builds on a **job scheduler**
 - ▶ for Hadoop: yarn
- ▶ considers **data is partitioned** over nodes
 - ▶ for Hadoop: blocks of a file in a distributed filesystem
- ▶ high level abstraction
 - ▶ programmer only specifies a **map** and a **reduce** function

Map-Reduce / Idea

- ▶ represent input and output data as key/value pairs.
- ▶ break down computation into three phases:
 1. **map phase**
 - ▶ apply a function **map** to each input key/value pair
 - ▶ the result is represented also as key/value pairs
 - ▶ execute map on each data node for its data partition (**data locality**)
 2. **shuffle phase**
 - ▶ group all intermediate key/value pairs by key to key/valueset pairs
 - ▶ repartition the intermediate data by key
 3. **reduce phase**
 - ▶ apply a function **reduce** to each intermediate key/valueset pair
 - ▶ execute reduce on each node for its intermediate data partition

Note: The **shuffle phase** is also called **sort** or **merge** phase.

Map-Reduce



Key-Value input data

- ▶ Map-Reduce requires the data to be stored in a key-value format
- ▶ Examples:

Key	Value
document	array of words
document	word
user	movies
user	friends
user	tweet

The Paradigm - Formally

Given

- ▶ A set of input keys I
- ▶ A set of output keys O
- ▶ A set of input values X
- ▶ A set of intermediate values V
- ▶ A set of output values Y

We can define:

$$\text{map} : I \times X \rightarrow (O \times V)^*$$

and

$$\text{reduce} : O \times (V^*) \rightarrow O \times Y$$

where $*$ denotes sequences.

Word Count Example

Map:

- ▶ Input: document-word list pairs
- ▶ Output: word-count pairs


$$(d_n, (w_1, \dots, w_M)) \mapsto ((w, c_w)_{w \in W: c_w > 0})$$

Reduce:

- ▶ Input: word-(count list) pairs
- ▶ Output: word-count pairs

$$(w, (c_1, \dots, c_K)) \mapsto (w, \sum_{k=1}^K c_k)$$

Word Count Example



Hadoop Example / Map

```
1 public static class Map
2     extends MapReduceBase
3     implements Mapper<LongWritable, Text, Text, IntWritable> {
4     private final static IntWritable one = new IntWritable(1);
5     private Text word = new Text();
6
7     public void map(LongWritable key, Text value,
8                     OutputCollector<Text, IntWritable> output,
9                     Reporter reporter)
10        throws IOException {
11
12        String line = value.toString ();
13        StringTokenizer tokenizer = new StringTokenizer(line );
14
15        while ( tokenizer.hasMoreTokens() ) {
16            word.set( tokenizer.nextToken());
17            output.collect( word, one);
18        }
19    }
20 }
```

Hadoop Example / Reduce

```
1 public static class Reduce
2     extends MapReduceBase
3     implements Reducer<Text, IntWritable, Text, IntWritable> {
4
5     public void reduce(Text key, Iterator<IntWritable> values,
6                       OutputCollector<Text, IntWritable> output,
7                       Reporter reporter)
8         throws IOException {
9
10        int sum = 0;
11        while (values.hasNext())
12            sum += values.next().get();
13
14        output.collect(key, new IntWritable(sum));
15    }
16 }
```

Hadoop Example / Main

```
1 public static void main(String[] args) throws Exception {
2     JobConf conf = new JobConf(WordCount.class);
3     conf.setJobName("wordcount");
4
5     conf.setOutputKeyClass(Text.class);
6     conf.setOutputValueClass(IntWritable.class);
7
8     conf.setMapperClass(Map.class);
9     conf.setCombinerClass(Reduce.class);
10    conf.setReducerClass(Reduce.class);
11
12    conf.setInputFormat(TextInputFormat.class);
13    conf.setOutputFormat(TextOutputFormat.class);
14
15    FileInputFormat.setInputPaths(conf, new Path(args[0]));
16    FileOutputFormat.setOutputPath(conf, new Path(args[1]));
17
18    JobClient.runJob(conf);
19 }
20 }
```


Execution

- ▶ mappers are executed in parallel
- ▶ reducers are executed in parallel
 - ▶ started after all mappers have completed
- ▶ high-level abstraction:
 - ▶ No need to worry about how many processors are available
 - ▶ No need to specify which ones will be mappers and which ones will be reducers

Fault Tolerance

- ▶ **map failure**
 - ▶ re-execute map
 - ▶ preferably on another node
 - ▶ speculative execution
 - ▶ execute two mappers on each data segment in parallel each
 - ▶ keep results from the first
 - ▶ kill the slower one, once the other completed.
- ▶ **node failure**
 - ▶ re-execute completed and in-progress map()
 - ▶ re-execute in-progress reduce tasks
- ▶ **particular key-value pairs** that cause mappers to crash
 - ▶ skip just the problematic pairs

Parallel Efficiency of Map-Reduce

- ▶ We have p processors for performing *map* and *reduce* operations
- ▶ Time to perform a task T on data D : $t(T, 1) = wD$
- ▶ Time for producing intermediate data σD after the *map* phase:
 $t(T^{\text{inter}}, 1) = \sigma D$
- ▶ Overheads:
 - ▶ intermediate data per mapper: $\frac{\sigma D}{p}$
 - ▶ each of the p reducers needs to read one p -th of the data written by each of the p mappers:

$$\frac{\sigma D}{p} \frac{1}{p} p = \frac{\sigma D}{p}$$

- ▶ Time for performing the task with Map-reduce:

$$t_{MR}(T, p) = \frac{wD}{p} + 2K \frac{\sigma D}{p}$$

Note: K represents the overhead of IO operations (reading and writing data to disk)

Parallel Efficiency of Map-Reduce

- ▶ Time for performing the task in one processor: wD
- ▶ Time for performing the task with p processors on Map-reduce:

$$t_{MR}(T, p) = \frac{wD}{p} + 2K \frac{\sigma D}{p}$$

- ▶ Efficiency of Map-Reduce:

$$\begin{aligned}
 e_{MR}(T, p) &= \frac{t(T, 1)}{p \cdot t(T, p)} \\
 &= \frac{wD}{p \left(\frac{wD}{p} + 2K \frac{\sigma D}{p} \right)} \\
 &= \frac{1}{1 + 2K \frac{\sigma}{w}}
 \end{aligned}$$

Parallel Efficiency of Map-Reduce

$$e_{MR}(T, p) = \frac{1}{1 + 2K \frac{\sigma}{w}}$$

- ▶ Apparently the efficiency is independent of p
- ▶ High speedups can be achieved with large number of processors
- ▶ If σ is high (too much intermediate data) the efficiency deteriorates
- ▶ In many cases σ depends on p

Outline

1. Introduction
2. Parallel Computing Speedup
3. Example: Counting Words
4. Map-Reduce
5. Map-Reduce Tutorial

Mappers

- ▶ extend baseclass **Mapper<KI,VI,KO,VO>**
(org.apache.hadoop.mapreduce)
 - ▶ types: KI = input key, VI = input value,
KO = output key, VO = output value.

Mappers

- ▶ extend baseclass **Mapper<KI,VI,KO,VO>**
(org.apache.hadoop.mapreduce)
 - ▶ types: KI = input key, VI = input value,
KO = output key, VO = output value.

- ▶ overwrite **map(KI key, VI value, Context ctxt)**
 - ▶ **Context** is an inner class of **Mapper<KI,VI,KO,VO>**
 - ▶ **write(KO,VO)**: write next output pair

Mappers

- ▶ extend baseclass **Mapper<KI,VI,KO,VO>**
(org.apache.hadoop.mapreduce)
 - ▶ types: KI = input key, VI = input value,
KO = output key, VO = output value.
- ▶ overwrite **map(KI key, VI value, Context ctxt)**
 - ▶ **Context** is an inner class of **Mapper<KI,VI,KO,VO>**
 - ▶ **write(KO,VO)**: write next output pair
- ▶ optionally,
setup(Context ctxt): set up mapper
 - ▶ called once before first call to **map**
- ▶ **cleanup(Context ctxt)**: clean up mapper
 - ▶ called once after last call to **map**

Key and Value Types

Requirements for key and value types:

- ▶ serializable — **Writable** (org.apache.hadoop.io)
 - ▶ To store the output of mappers, combiners and reducers in files

Key and Value Types

Requirements for key and value types:

- ▶ serializable — **Writable** (org.apache.hadoop.io)
 - ▶ To store the output of mappers, combiners and reducers in files

Additional requirements for key types:

- ▶ comparable — **Comparable** (java.lang)
 - ▶ To sort the output of mappers and combiners by key

Key and Value Types

Requirements for key and value types:

- ▶ serializable — **Writable** (org.apache.hadoop.io)
 - ▶ To store the output of mappers, combiners and reducers in files

Additional requirements for key types:

- ▶ comparable — **Comparable** (java.lang)
 - ▶ To sort the output of mappers and combiners by key
- ▶ stable **hashCode** across different JVM instances
 - ▶ To partition the output of combiners by key.
 - ▶ The default implementation in **Object** is not stable!

Key and Value Types

Requirements for key and value types:

- ▶ serializable — **Writable** (org.apache.hadoop.io)
 - ▶ To store the output of mappers, combiners and reducers in files

Additional requirements for key types:

- ▶ comparable — **Comparable** (java.lang)
 - ▶ To sort the output of mappers and combiners by key
- ▶ stable **hashCode** across different JVM instances
 - ▶ To partition the output of combiners by key.
 - ▶ The default implementation in **Object** is not stable!
- ▶ pooled in interface **WritableComparable** (org.apache.hadoop.io)

Serialization

To store the output of mappers, combiners and reducers in files, keys and values have to be **serialized**:

- ▶ hadoop requires all keys and values of a step to be of the same class.
↪ the class of an object to deserialize is known in advance.
- ▶ standard Java serialization ([java.lang.Serializable](#)) serializes class information, thus is more verbose and complex and therefore not used in hadoop.
- ▶ new interface [Writable](#) (org.apache.hadoop.io):
 - ▶ [void write\(DataOutput out\) throws IOException](#):
write object to a data output.
 - ▶ [void readFields\(DataInput in\) throws IOException](#):
set member variables (fields) of an object to values read from a data input.
 - ▶ standard [DataInput](#) and [DataOutput](#) (java.io)

Serialization (2/2)

- ▶ **Writable** wrappers for elementary data types:

BooleanWritable	boolean
ByteWritable	byte
DoubleWritable	double
FloatWritable	float
IntWritable	int
LongWritable	long
ShortWritable	short
Text	String

(all in org.apache.hadoop.io)

- ▶ these are not subclasses of the default wrappers **Integer**, **Double** etc. (as the latter are final)

Serialization (2/2)

- ▶ **Writable** wrappers for elementary data types:

BooleanWritable	boolean
ByteWritable	byte
DoubleWritable	double
FloatWritable	float
IntWritable	int
LongWritable	long
ShortWritable	short
Text	String

(all in org.apache.hadoop.io)

- ▶ these are not subclasses of the default wrappers **Integer**, **Double** etc. (as the latter are final)
- ▶ If one needs to pass more complex objects between steps, implement custom **Writable** (in terms of these elementary **Writables**).

Example 1 / Mapper (in principle)

```
1 import java.io.IOException;
2 import java.util.StringTokenizer;
3
4 import org.apache.hadoop.io. IntWritable ;
5 import org.apache.hadoop.io. Text;
6 import org.apache.hadoop.mapreduce.Mapper;
7
8 public class TokenizerMapperS
9     extends Mapper<Object, Text, Text, IntWritable >{
10
11     public void map(Object key, Text value, Context context
12         ) throws IOException, InterruptedException {
13         StringTokenizer itr = new StringTokenizer(value.toString ());
14         while ( itr.hasMoreTokens())
15             context.write(new Text(itr.nextToken()), new IntWritable (1));
16     }
17 }
```

Example 1 / Mapper

```
1 import java.io.IOException;
2 import java.util.StringTokenizer;
3
4 import org.apache.hadoop.io. IntWritable ;
5 import org.apache.hadoop.io. Text;
6 import org.apache.hadoop.mapreduce. Mapper;
7
8 public class TokenizerMapper
9     extends Mapper<Object, Text, Text, IntWritable>{
10
11     private final static IntWritable one = new IntWritable(1);
12     private Text word = new Text();
13
14     public void map(Object key, Text value, Context context
15         ) throws IOException, InterruptedException {
16         StringTokenizer itr = new StringTokenizer(value.toString());
17         while (itr.hasMoreTokens()) {
18             word.set(itr.nextToken());
19             context.write(word, one);
20         }
21     }
22 }
```

Job Configuration

Configuration (org.apache.hadoop.conf):

- ▶ default constructor:
read default configuration from files
 - ▶ core-default.xml and
 - ▶ core-site.xml(to be found in the classpath).
- ▶ **addResource(Path file):**
update configuration from another configuration file.
- ▶ **String get(String name):**
get the value of a configuration option.
- ▶ **set(String name, String value):**
set the value of a configuration option.

Job

Job (org.apache.hadoop.conf):

- ▶ constructor **Job(Configuration conf, String name)**:
create a new job.
- ▶ **setMapperClass(Class cls), setCombinerClass(Class cls), setReducerClass(Class cls)**:
set the class for mappers, combiners and reducers
- ▶ **setOutputKeyClass(Class cls), setOutputValueClass(Class cls)**:
set the class for output keys and values.
- ▶ **boolean waitForCompletion(boolean verbose)**:
submit job and wait until it completes.

Input and Output Paths

FileInputFormat (org.apache.hadoop.mapreduce.lib.input):

- ▶ **addInputPath(Job job, Path path):**
add input paths

FileOutputFormat (org.apache.hadoop.mapreduce.lib.output):

- ▶ **setOutputPath(Job job, Path path):**
set output paths

Example 1 / Job Runner (Mapper Only)

```
1 import java.io.IOException;
2 import java.util.StringTokenizer;
3
4 import org.apache.hadoop.conf.Configuration;
5 import org.apache.hadoop.fs.Path;
6 import org.apache.hadoop.io.IntWritable;
7 import org.apache.hadoop.io.Text;
8 import org.apache.hadoop.mapreduce.Job;
9 import org.apache.hadoop.mapreduce.lib.input.FileInputFormat;
10 import org.apache.hadoop.mapreduce.lib.output.FileOutputFormat;
11
12 public class MRJobStarter1 {
13     public static void main(String[] args) throws Exception {
14         Configuration conf = new Configuration();
15         Job job = Job.getInstance(conf, "word count");
16         job.setMapperClass(TokenMapper.class);
17         job.setOutputKeyClass(Text.class);
18         job.setOutputValueClass(IntWritable.class);
19         FileInputFormat.addInputPath(job, new Path(args[0]));
20         FileOutputFormat.setOutputPath(job, new Path(args[1]));
21
22         System.exit(job.waitForCompletion(true) ? 0 : 1);
23     }
24 }
```

Compiling and Running Map Reduce Classes

1. Set paths

```
1 export PATH=PATH : /home/lst/system/hadoop/binexportHADOOP_CLASSPATH=(JAVA_HOME)/lib/to
```

Compiling and Running Map Reduce Classes

1. Set paths

```
1 export PATH=PATH : /home/lst/system/hadoop/binexportHADOOP_CLASSPATH=(JAVA_HOME)/lib/to
```

2. Compile sources:

```
1 hadoop com.sun.tools.javac.Main --sourcepath . MRJobStarter1.java
```


Compiling and Running Map Reduce Classes

1. Set paths

```
1 export PATH=PATH : /home/lst/system/hadoop/binexportHADOOP_CLASSPATH=(JAVA_HOME)/lib/to
```

2. Compile sources:

```
1 hadoop com.sun.tools.javac.Main --sourcepath . MRJobStarter1.java
```

3. Package all class files of the job into a jar:

```
1 jar cf job.jar MRJobStarter1.class TokenizerMapper.class
```

Compiling and Running Map Reduce Classes

1. Set paths

```
1 export PATH=PATH : /home/lst/system/hadoop/binexportHADOOP_CLASSPATH=(JAVA_HOME)/lib/to
```

2. Compile sources:

```
1 hadoop com.sun.tools.javac.Main --sourcepath . MRJobStarter1.java
```

3. Package all class files of the job into a jar:

```
1 jar cf job.jar MRJobStarter1.class TokenizerMapper.class
```

4. Run the jar:

```
1 hadoop jar job.jar MRJobStarter1 /ex1/input /ex1/output
```

- ▶ the output directory must not yet exist.

Example 1 / Inputs and Output

▶ input 1:

```
1 Hello World Bye World
```

▶ input 2:

```
1 Hello Hadoop Goodbye Hadoop
```

▶ output:

```
lst@lst-uni:~> hdfs dfs -ls /ex1/output.ex2
Found 2 items
-rw-r-r-  2 lst supergroup          0 2016-05-24 18:59 /ex1/output.ex2/_SUCCESS
-rw-r-r-  2 lst supergroup        66 2016-05-24 18:59 /ex1/output.ex2/part-r-00000
lst@lst-uni:~> hdfs dfs -cat /ex1/output.ex2/part-r-00000
Bye      1
Goodbye 1
Hadoop  1
Hadoop  1
Hello   1
Hello   1
World   1
World   1
```

Reducers

- ▶ extend baseclass **Reducer**<KI,VI,KO,VO>
(org.apache.hadoop.mapreduce)
 - ▶ types: KI = input key, VI = input value,
KO = output key, VO = output value.

Reducers

- ▶ extend baseclass **Reducer<KI,VI,KO,VO>**
(org.apache.hadoop.mapreduce)
 - ▶ types: KI = input key, VI = input value,
KO = output key, VO = output value.

- ▶ overwrite **reduce(KI key, Iterable<VI> value, Context ctxt)**
 - ▶ **Context** is an inner class of **Reducer<KI,VI,KO,VO>**
 - ▶ **write(KO,VO)**: write next output pair

Reducers

- ▶ extend baseclass **Reducer<KI,VI,KO,VO>**
(org.apache.hadoop.mapreduce)
 - ▶ types: KI = input key, VI = input value,
KO = output key, VO = output value.
- ▶ overwrite **reduce(KI key, Iterable<VI> value, Context ctxt)**
 - ▶ **Context** is an inner class of **Reducer<KI,VI,KO,VO>**
 - ▶ **write(KO,VO)**: write next output pair
- ▶ optionally,
setup(Context ctxt): set up reducer
 - ▶ called once before first call to **reduce**
- ▶ **cleanup(Context ctxt)**: clean up reducer
 - ▶ called once after last call to **reduce**

Example 2 / Reducer

```
1 import java.io.IOException;
2
3 import org.apache.hadoop.io. IntWritable ;
4 import org.apache.hadoop.io. Text;
5 import org.apache.hadoop.mapreduce.Reducer;
6
7 public class IntSumReducer
8     extends Reducer<Text,IntWritable,Text, IntWritable > {
9     private IntWritable result = new IntWritable();
10
11     public void reduce(Text key, Iterable<IntWritable> values,
12                       Context context
13                       ) throws IOException, InterruptedException {
14         int sum = 0;
15         for ( IntWritable val : values)
16             sum += val.get();
17         result.set(sum);
18         context.write(key, result );
19     }
20 }
```

Example 2 / Job Runner

```
1 import java.io.IOException;
2 import java.util.StringTokenizer;
3
4 import org.apache.hadoop.conf. Configuration ;
5 import org.apache.hadoop.fs. Path;
6 import org.apache.hadoop.io. IntWritable ;
7 import org.apache.hadoop.io. Text;
8 import org.apache.hadoop.mapreduce. Job;
9 import org.apache.hadoop.mapreduce.lib.input. FileInputFormat ;
10 import org.apache.hadoop.mapreduce.lib.output. FileOutputFormat;
11
12 public class MRJobStarter2 {
13     public static void main(String[] args) throws Exception {
14         Configuration conf = new Configuration ();
15         Job job = Job.getInstance (conf, "word count");
16         job.setMapperClass (TokenizerMapper.class);
17         job.setReducerClass (IntSumReducer.class);
18         job.setOutputKeyClass (Text.class);
19         job.setOutputValueClass (IntWritable.class);
20         FileInputFormat.addInputPath (job, new Path (args [0]));
21         FileOutputFormat.setOutputPath (job, new Path (args [1]));
22
23         System.exit (job.waitForCompletion (true) ? 0 : 1);
24     }
25 }
```


Example 2 / Output

```
lst@lst-uni:~> hdfs dfs -cat /ex1/output.2/part*  
Bye 1  
Goodbye 1  
Hadoop 2  
Hello 2  
World 2
```

Example 3 / Job Runner: Only Mapper and Combiner

```
1 import java.io.IOException;
2 import java.util.StringTokenizer;
3
4 import org.apache.hadoop.conf. Configuration ;
5 import org.apache.hadoop.fs. Path;
6 import org.apache.hadoop.io. IntWritable ;
7 import org.apache.hadoop.io. Text;
8 import org.apache.hadoop.mapreduce. Job;
9 import org.apache.hadoop.mapreduce.lib.input. FileInputFormat ;
10 import org.apache.hadoop.mapreduce.lib.output. FileOutputFormat;
11
12 public class MRJobStarter3 {
13     public static void main(String[] args) throws Exception {
14         Configuration conf = new Configuration ();
15         Job job = Job.getInstance (conf, "word count");
16         job.setMapperClass (TokenizerMapper.class);
17         job.setCombinerClass (IntSumReducer.class);
18         job.setOutputKeyClass (Text.class);
19         job.setOutputValueClass (IntWritable.class);
20         FileInputFormat.addInputPath (job, new Path (args [0]));
21         FileOutputFormat.setOutputPath (job, new Path (args [1]));
22
23         System.exit (job.waitForCompletion (true) ? 0 : 1);
24     }
25 }
```

Example 3 / Output

```
lst@lst-uni:~> hdfs dfs -cat /ex1/output.3/part*
Bye      1
Goodbye  1
Hadoop   2
Hello    1
Hello    1
World    2
```

Example 4 / Job Runner: Combiner and Reducer

```
1 import java.io.IOException;
2 import java.util.StringTokenizer;
3
4 import org.apache.hadoop.conf. Configuration;
5 import org.apache.hadoop.fs. Path;
6 import org.apache.hadoop.io. IntWritable;
7 import org.apache.hadoop.io. Text;
8 import org.apache.hadoop.mapreduce. Job;
9 import org.apache.hadoop.mapreduce.lib.input. FileInputFormat;
10 import org.apache.hadoop.mapreduce.lib.output. FileOutputFormat;
11
12 public class MRJobStarter4 {
13     public static void main(String[] args) throws Exception {
14         Configuration conf = new Configuration();
15         Job job = Job.getInstance(conf, "word count");
16         job.setMapperClass(TokenizerMapper.class);
17         job.setCombinerClass(IntSumReducer.class);
18         job.setReducerClass(IntSumReducer.class);
19         job.setOutputKeyClass(Text.class);
20         job.setOutputValueClass(IntWritable.class);
21         FileInputFormat.addInputPath(job, new Path(args[0]));
22         FileOutputFormat.setOutputPath(job, new Path(args[1]));
23
24         System.exit(job.waitForCompletion(true) ? 0 : 1);
25     }
26 }
```

The output is the same as only with a reducer,
but less intermediate data is moved.

Predefined Mappers

- ▶ **Mapper**: (org.apache.hadoop.mapreduce):

$$m(k, v) := ((k, v))$$

- ▶ **InverseMapper** (org.apache.hadoop.mapreduce.lib.map):

$$m(k, v) := ((v, k))$$

- ▶ **ChainMapper** (org.apache.hadoop.mapreduce.lib.chain)

$$\text{chain}_{m,\ell}(k, v) := (\ell(k', m') \mid (k', v') \in m(k, v))$$

- ▶ **TokenCounterMapper** (org.apache.hadoop.mapreduce.lib.map):

$$m(k, v) := ((k', 1) \mid k' \in \text{tokenize}(v))$$

Predefined Mappers (2/2)

- ▶ **RegexMapper** (`org.apache.hadoop.mapreduce.lib.map`):

$$m(k, v) := ((k', 1) \mid k' \in \text{find-regex}(v))$$

Regex pattern to search for and
groups to report

can be set via configuration options

PATTERN `mapreduce.mapper.regex`

GROUP `mapreduce.mapper.regexmapper..group`

Predefined Reducers

- ▶ **Reducer**: (`org.apache.hadoop.mapreduce`):

$$r(k, V) := (k, V)$$

- ▶ **IntSumReducer**, **LongSumReducer**
(`org.apache.hadoop.mapreduce.lib.reduce`):

$$m(k, V) := (v, \sum_{v \in V} v)$$

- ▶ **ChainReducer** (`org.apache.hadoop.mapreduce.lib.chain`)

$$\text{chain}_{r,s}(k, V) := s(r(k, V))$$

Job Default Values

```
mapper class: class org.apache.hadoop.mapreduce.Mapper
combiner class: null
reducer class: class org.apache.hadoop.mapreduce.Reducer
output key class: class org.apache.hadoop.io.LongWritable
output value class: class org.apache.hadoop.io.Text
```


Further Topics

- ▶ Controlling map-reduce jobs
 - ▶ YARN
 - ▶ chaining map-reduce jobs
 - ▶ iterative algorithms
- ▶ Controlling the number of mappers and reducers
- ▶ Managing resources required by all mappers or reducers
- ▶ Input and output using relational databases
- ▶ Streaming
- ▶ Examples, examples, examples

Summary

- ▶ Map Reduce is a **distributed computing framework**.
- ▶ Map Reduce represents input and output **data as key/value pairs**.
- ▶ Map Reduce decomposes computation into three phases:
 1. **map**: applying a function to **each input key/value pair**.
 2. **shuffle**:
 - ▶ **grouping intermediate data** into key/valueset pairs
 - ▶ **repartitioning intermediate data** by key over nodes
 3. **reduce**: applying a function to **each intermediate key/valueset pair**.
- ▶ Mappers are executed **data local**
- ▶ For a program, the **map** and **reduce** functions have to be specified.
 - ▶ the shuffle step is fixed.
- ▶ The size of the intermediate data is crucial for efficiency as it has to **be repartitioned**.

Further Readings

- ▶ original Map Reduce framework by Google:
 - ▶ Dean and Ghemawat [2004, 2008, 2010]
- ▶ MapReduce reference implementation in Hadoop:
 - ▶ [White, 2015, ch. 2, 7]
 - ▶ also ch. 6, 8 and 9.
- ▶ MapReduce in a document database, MongoDB:
 - ▶ [Chodorow, 2013, ch. 7]

References

- Kristina Chodorow. *MongoDB: The Definitive Guide*. O'Reilly and Associates, Beijing, 2 edition, May 2013. ISBN 978-1-4493-4468-9.
- Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. In *OSDI'04 Proceedings of the 6th Conference on Symposium on Operating Systems Design & Implementation*, volume 6, 2004.
- Jeffrey Dean and Sanjay Ghemawat. MapReduce. *Communications of the ACM*, 51(1):107, January 2008. ISSN 00010782. doi: 10.1145/1327452.1327492.
- Jeffrey Dean and Sanjay Ghemawat. MapReduce: A flexible data processing tool. *Communications of the ACM*, 53(1):72–77, 2010.
- Tom White. *Hadoop: The Definitive Guide*. O'Reilly, 4 edition, 2015.