

Big_{ML}

Some tools & techniques for large-scale ML

Joaquin Vanschoren



Universiteit Leiden

Big_{ML}

Some tools & techniques for large-scale ML

Joaquin Vanschoren

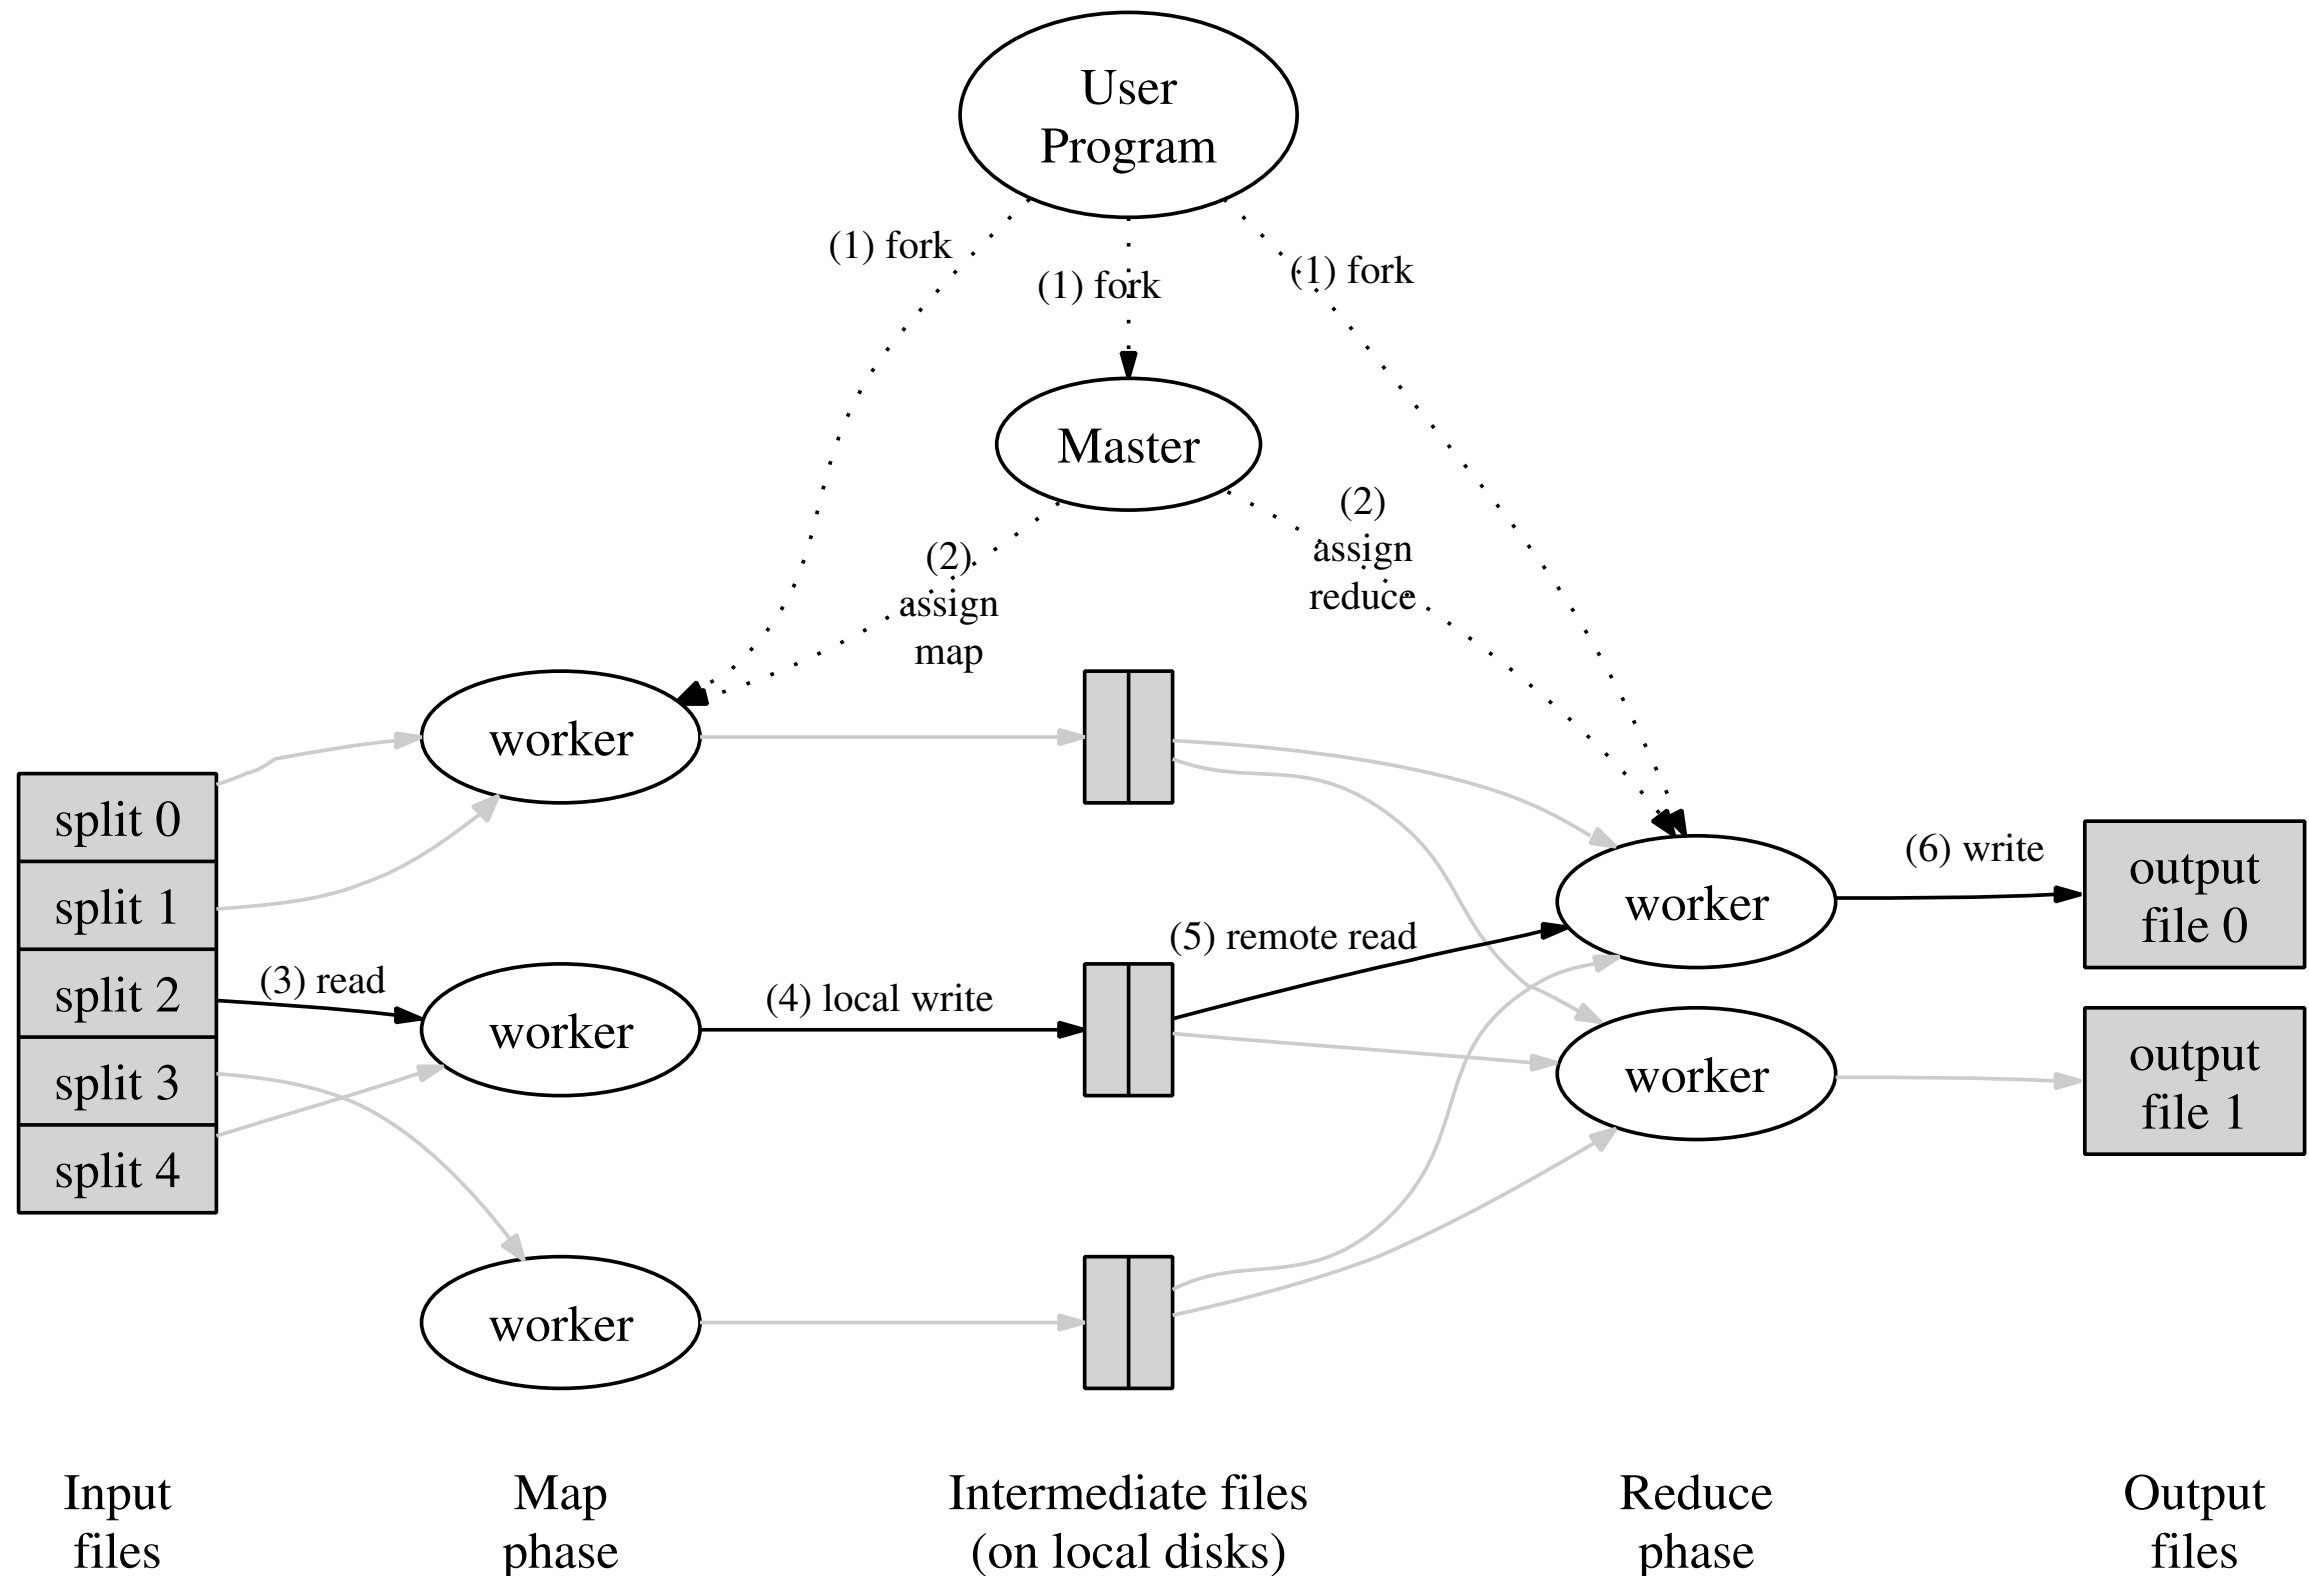
The Why

- A lot of ML problems concern really large datasets
 - Huge graphs: internet, social networks, protein interactions
 - Sensor data (cfr. Kristian)
 - InfraWatch: 5GB/day, 2TB/year, 8TB/4years @50MB/s -> 2 days
- Data is just getting bigger...

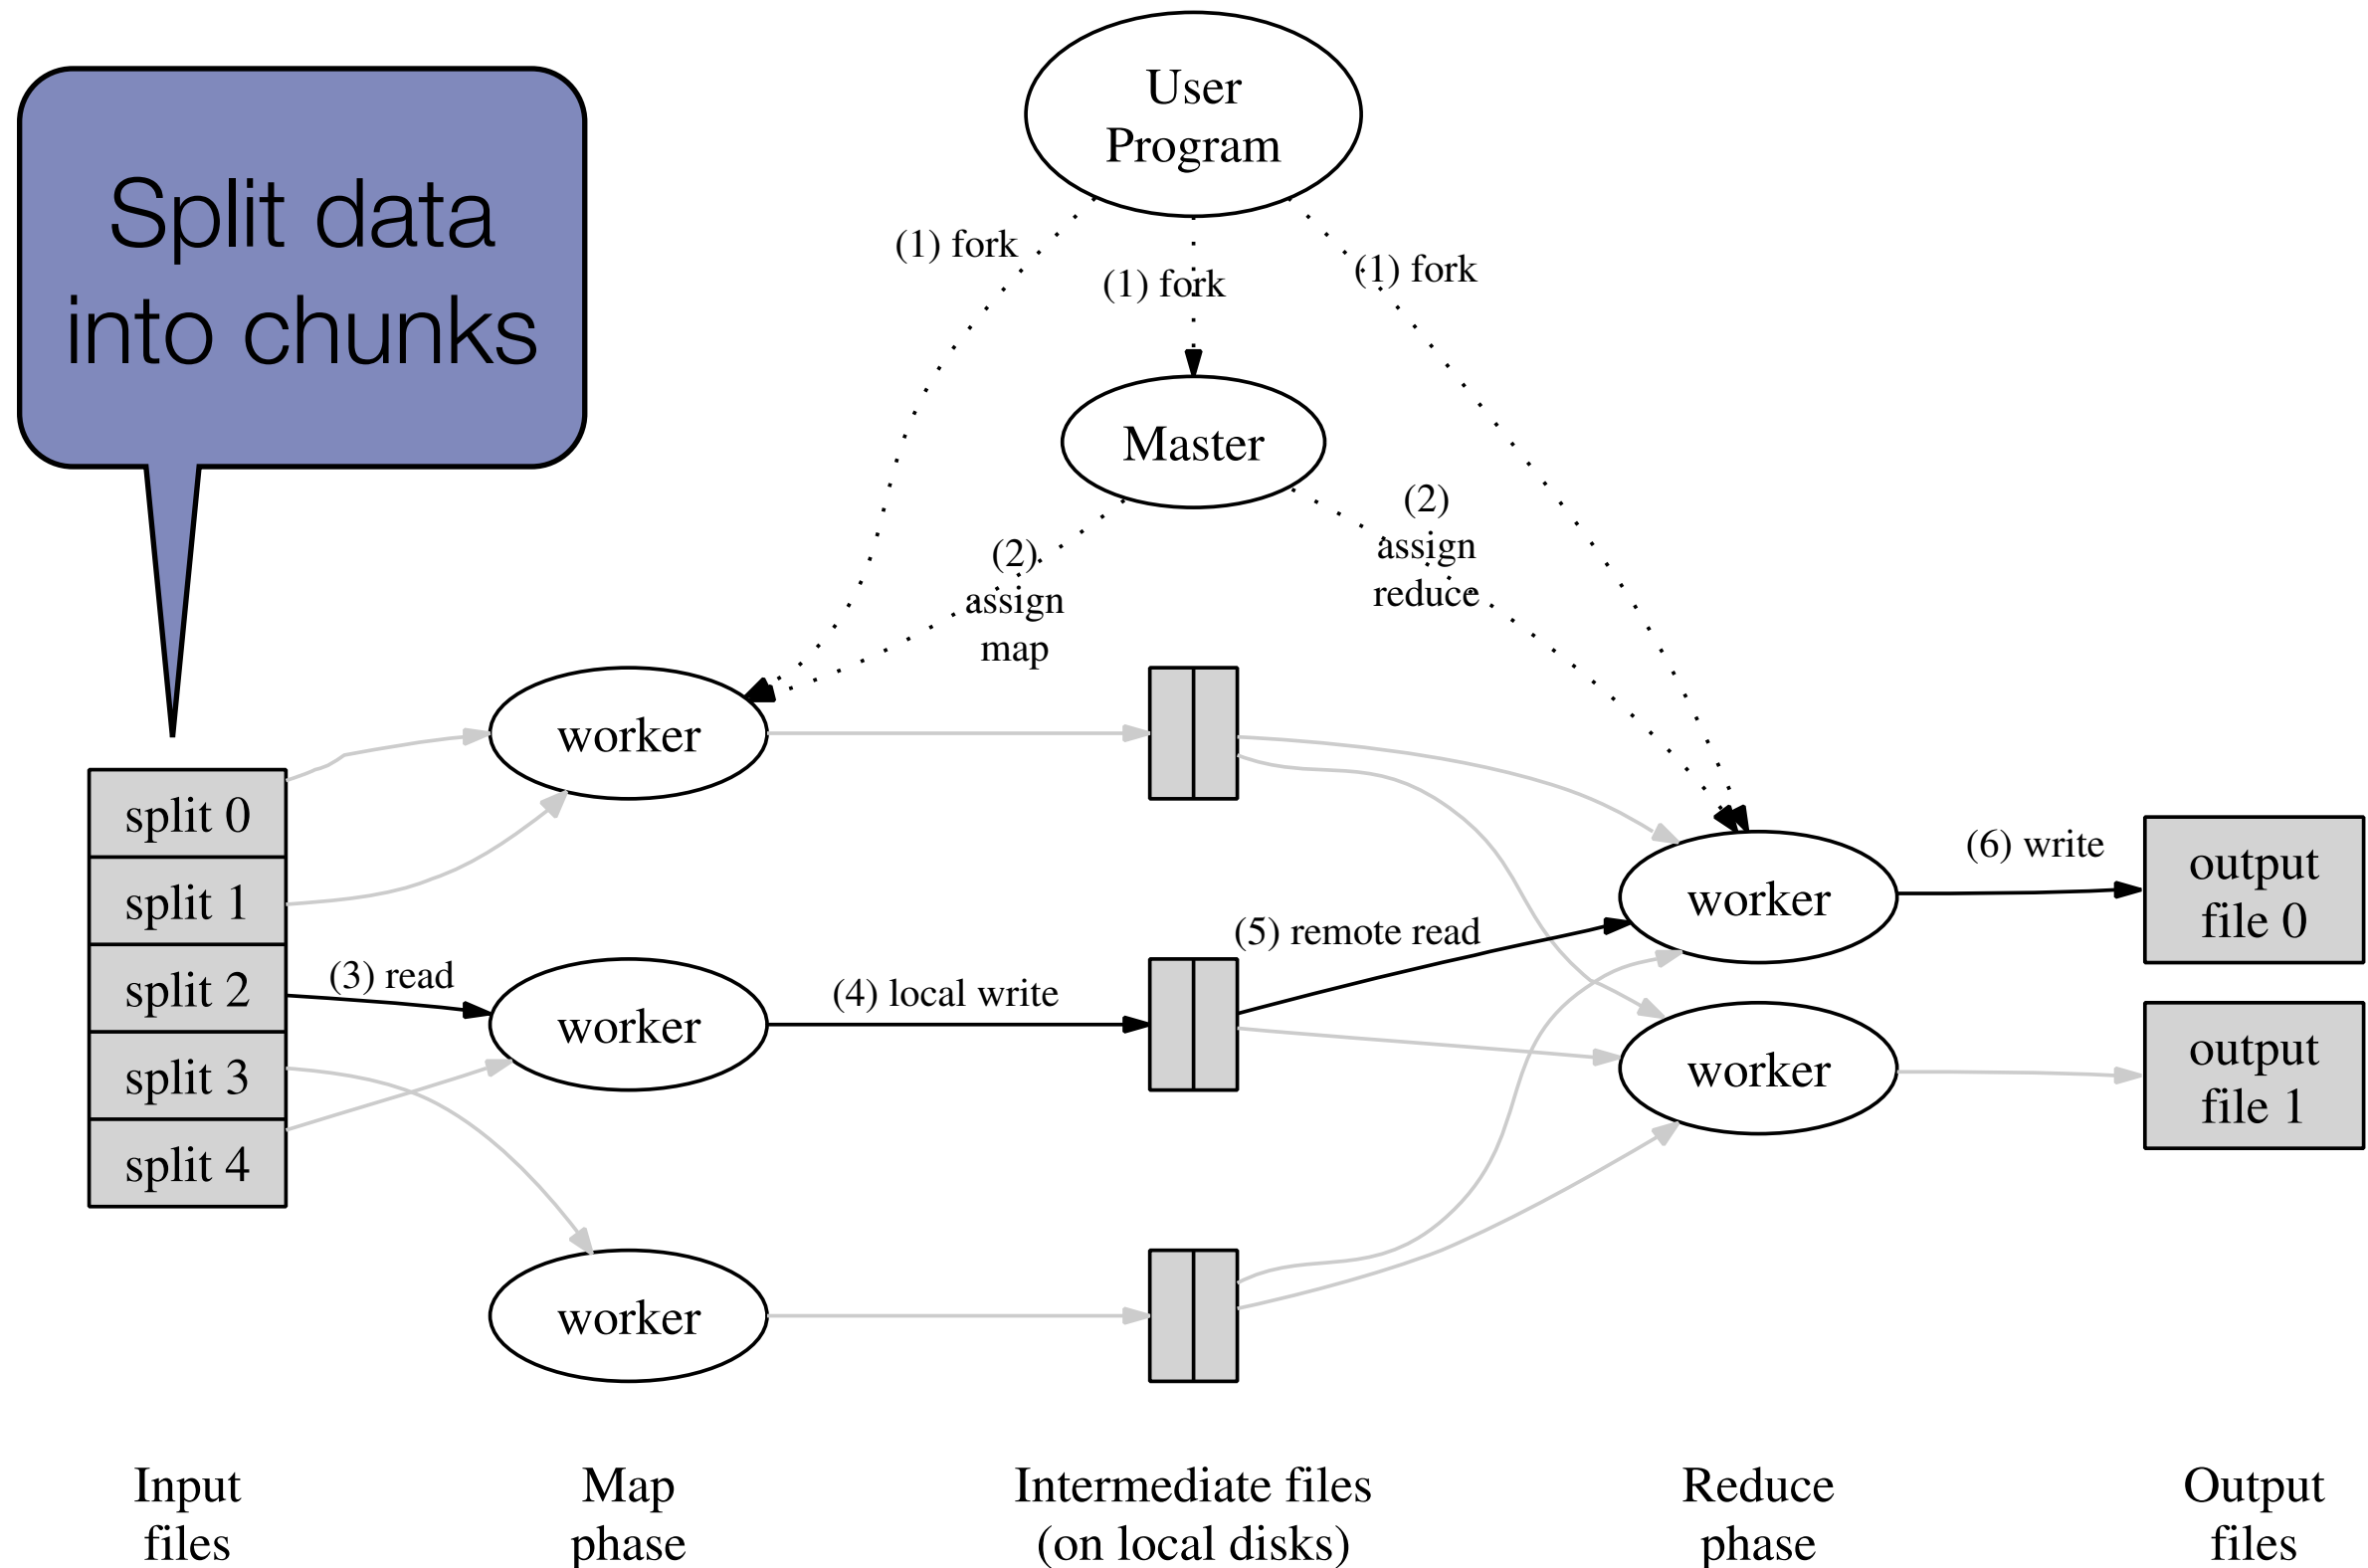
The What

- Map-Reduce (Hadoop) -> I/O bound processes
- Distributed computing -> CPU bound processes

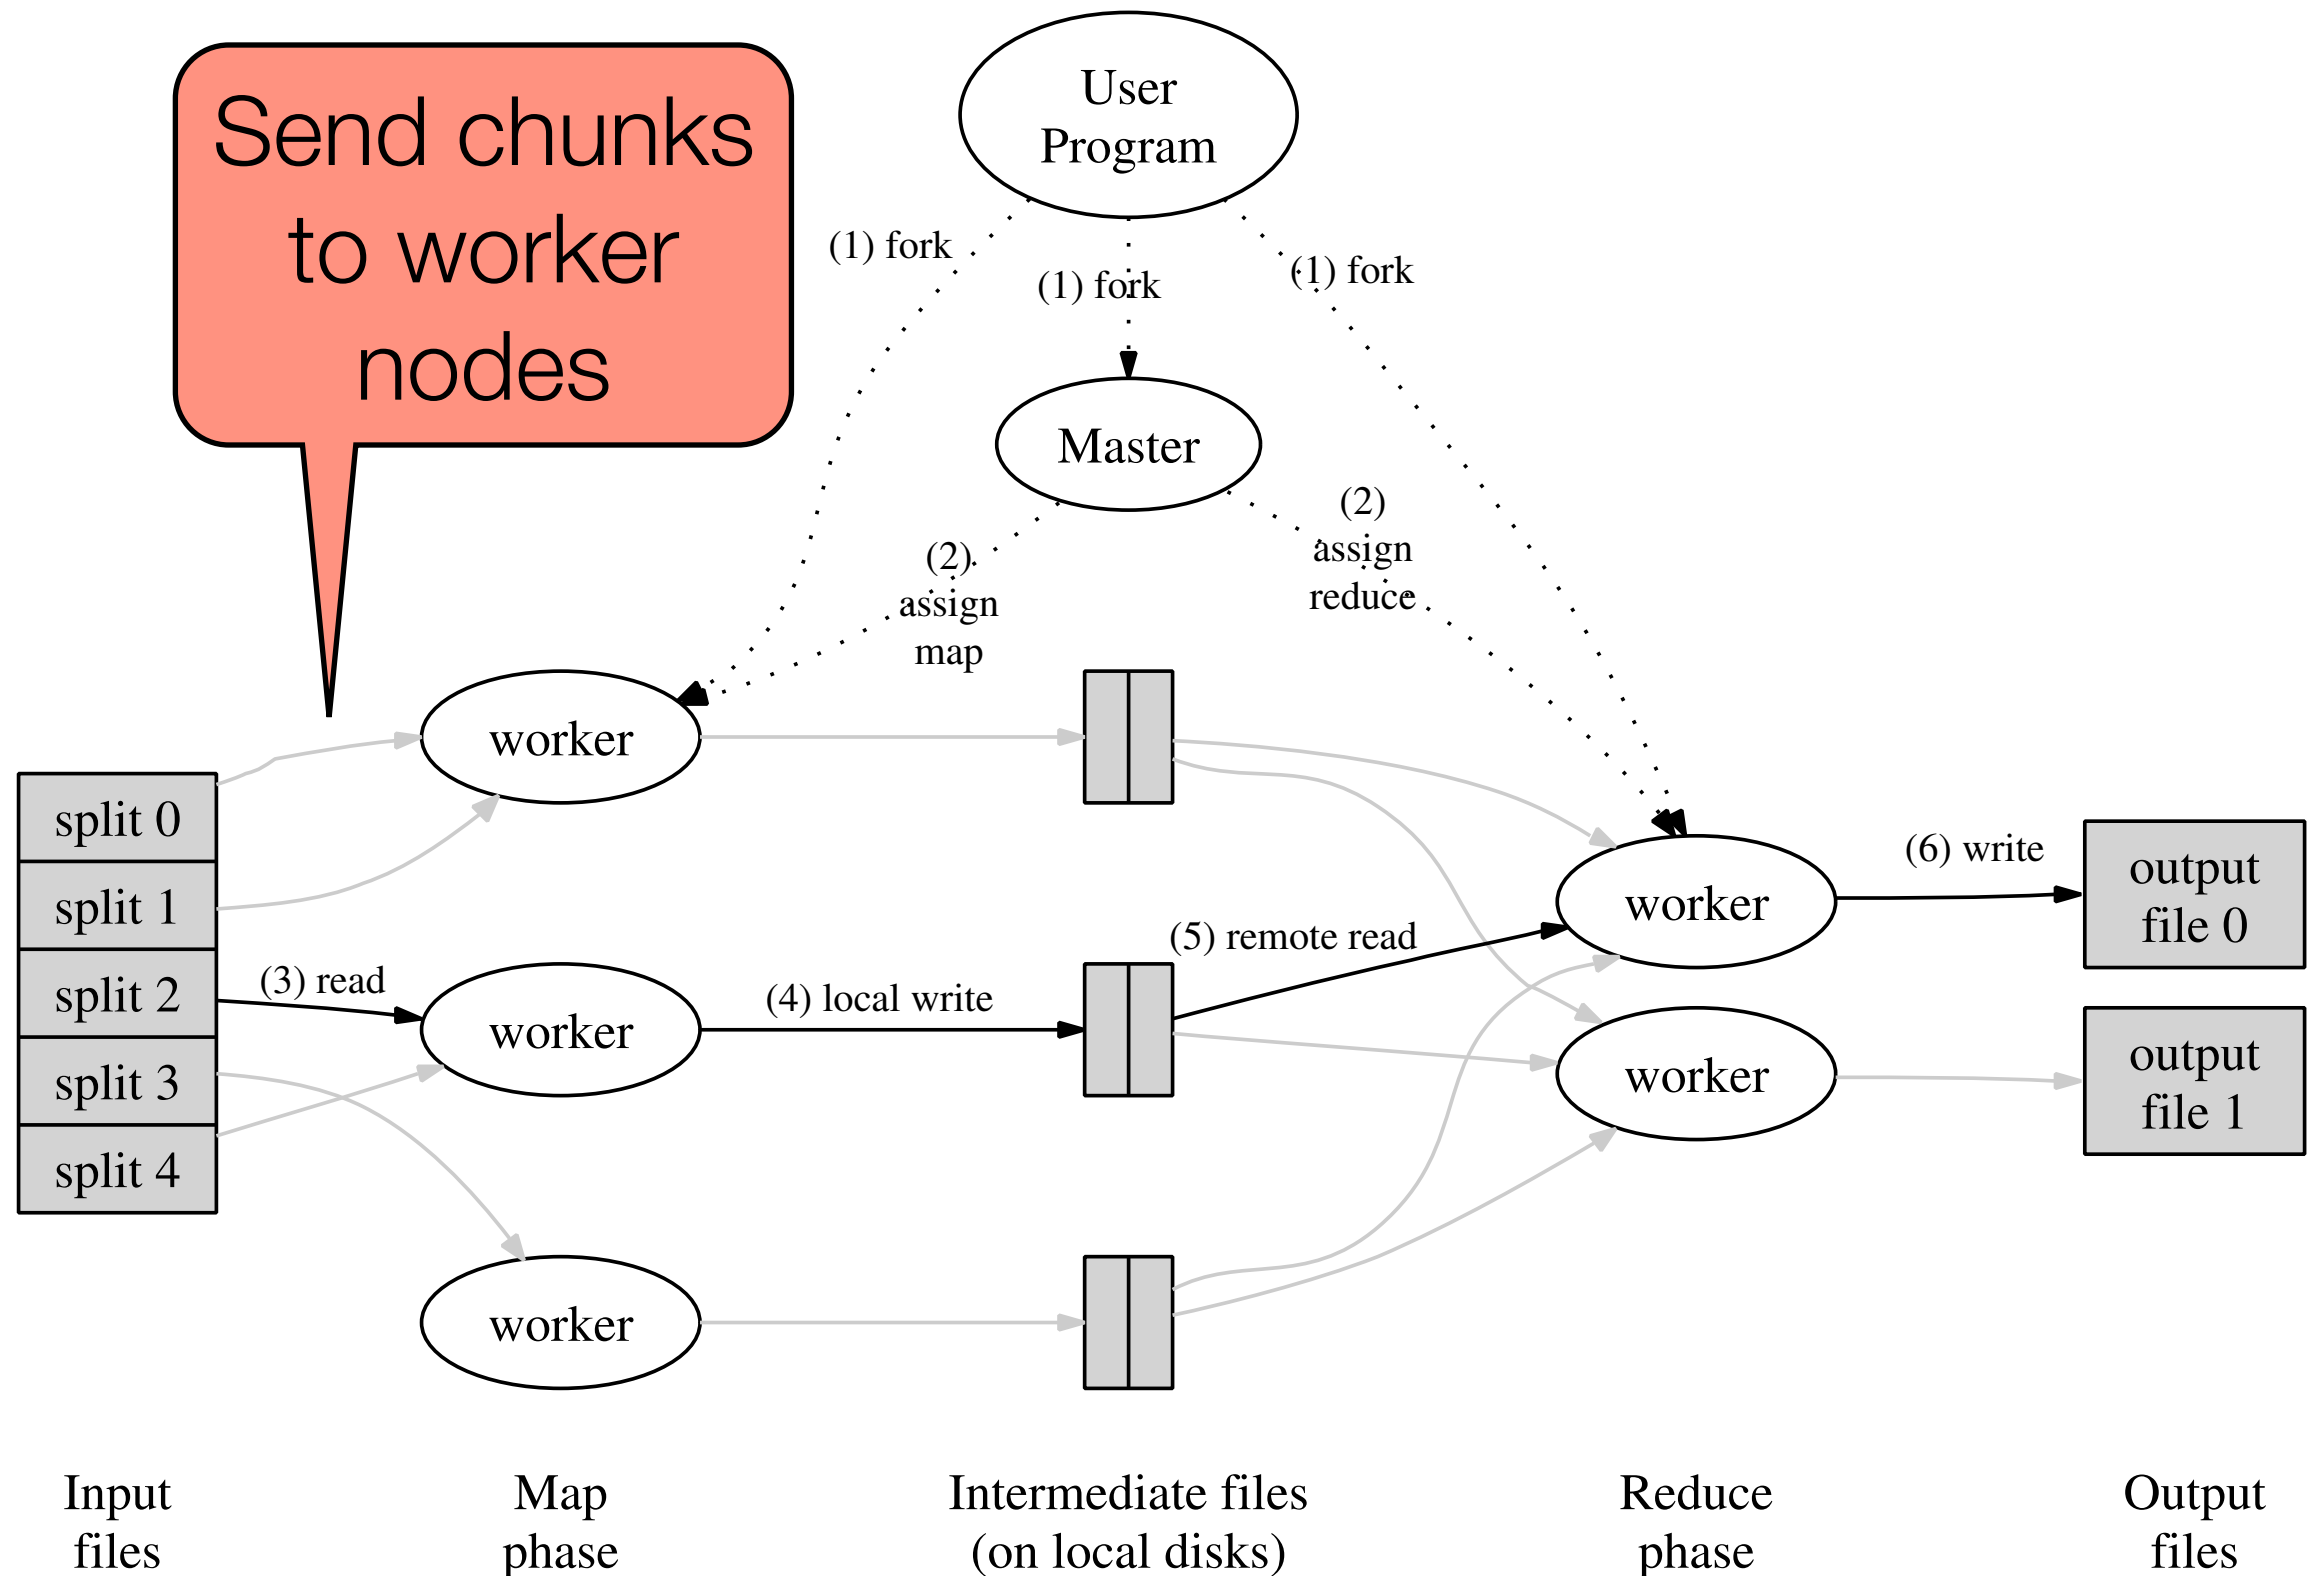
Map-reduce: distribute I/O over many nodes



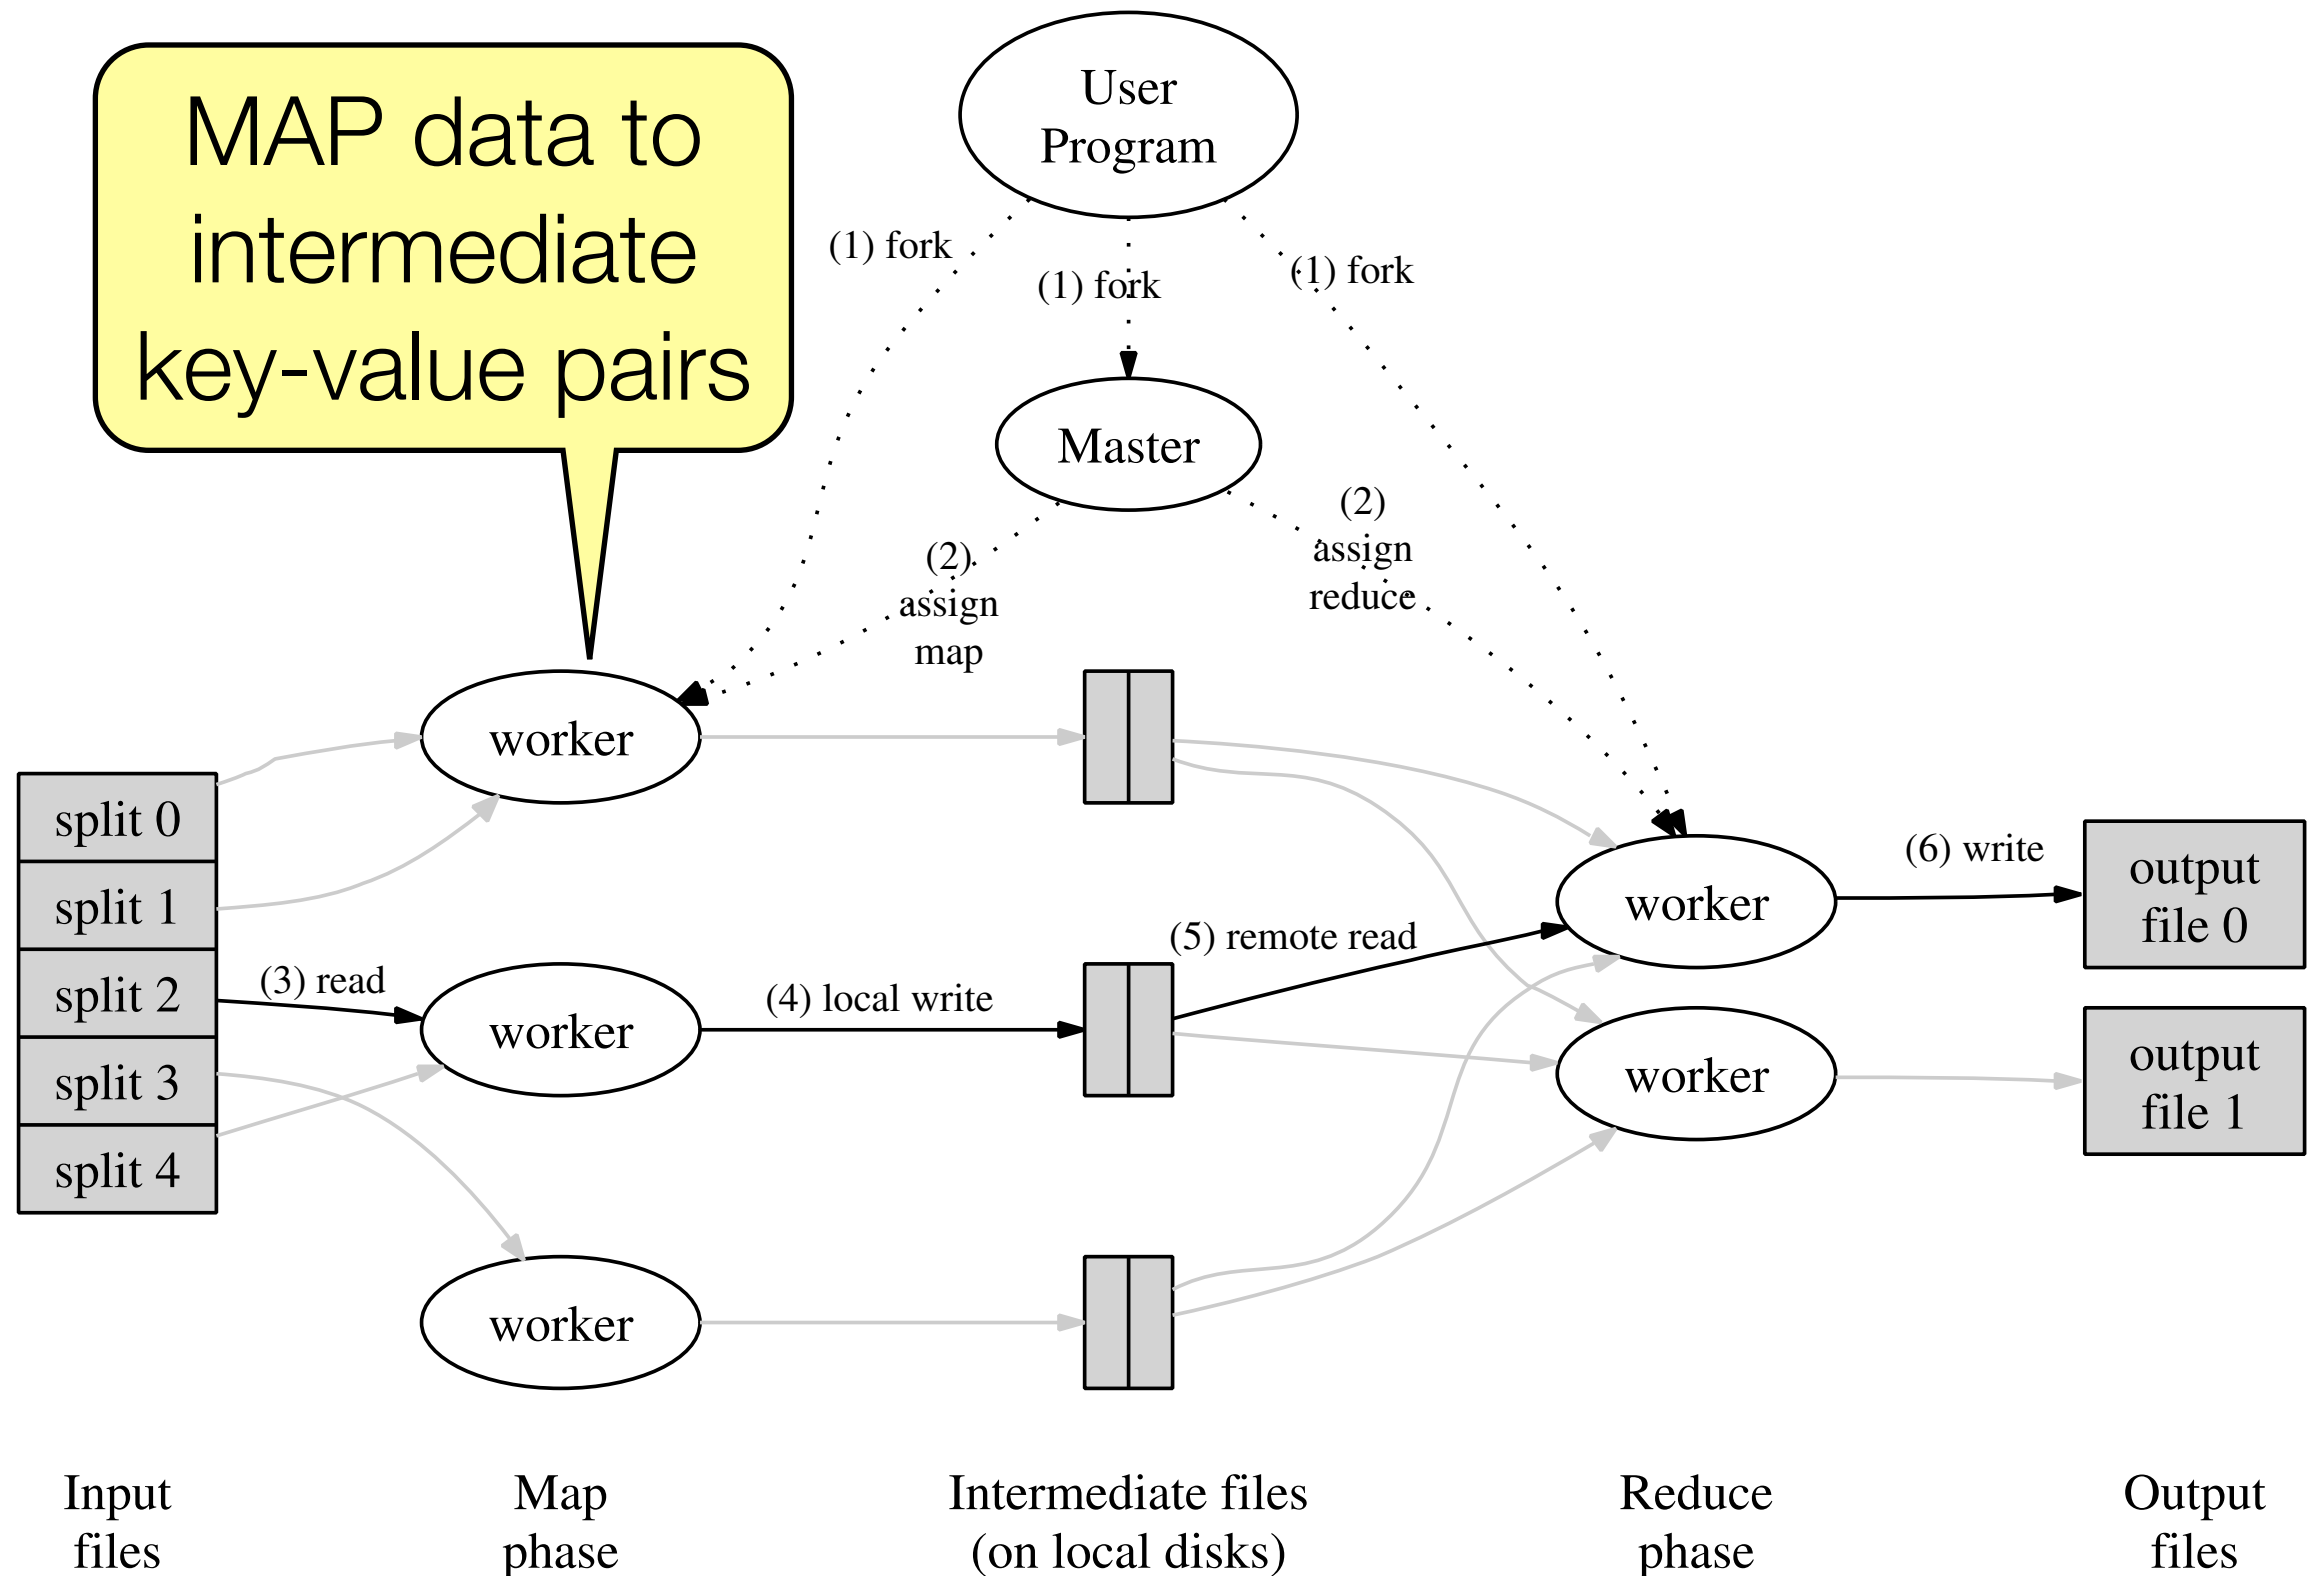
Map-reduce: distribute I/O over many nodes



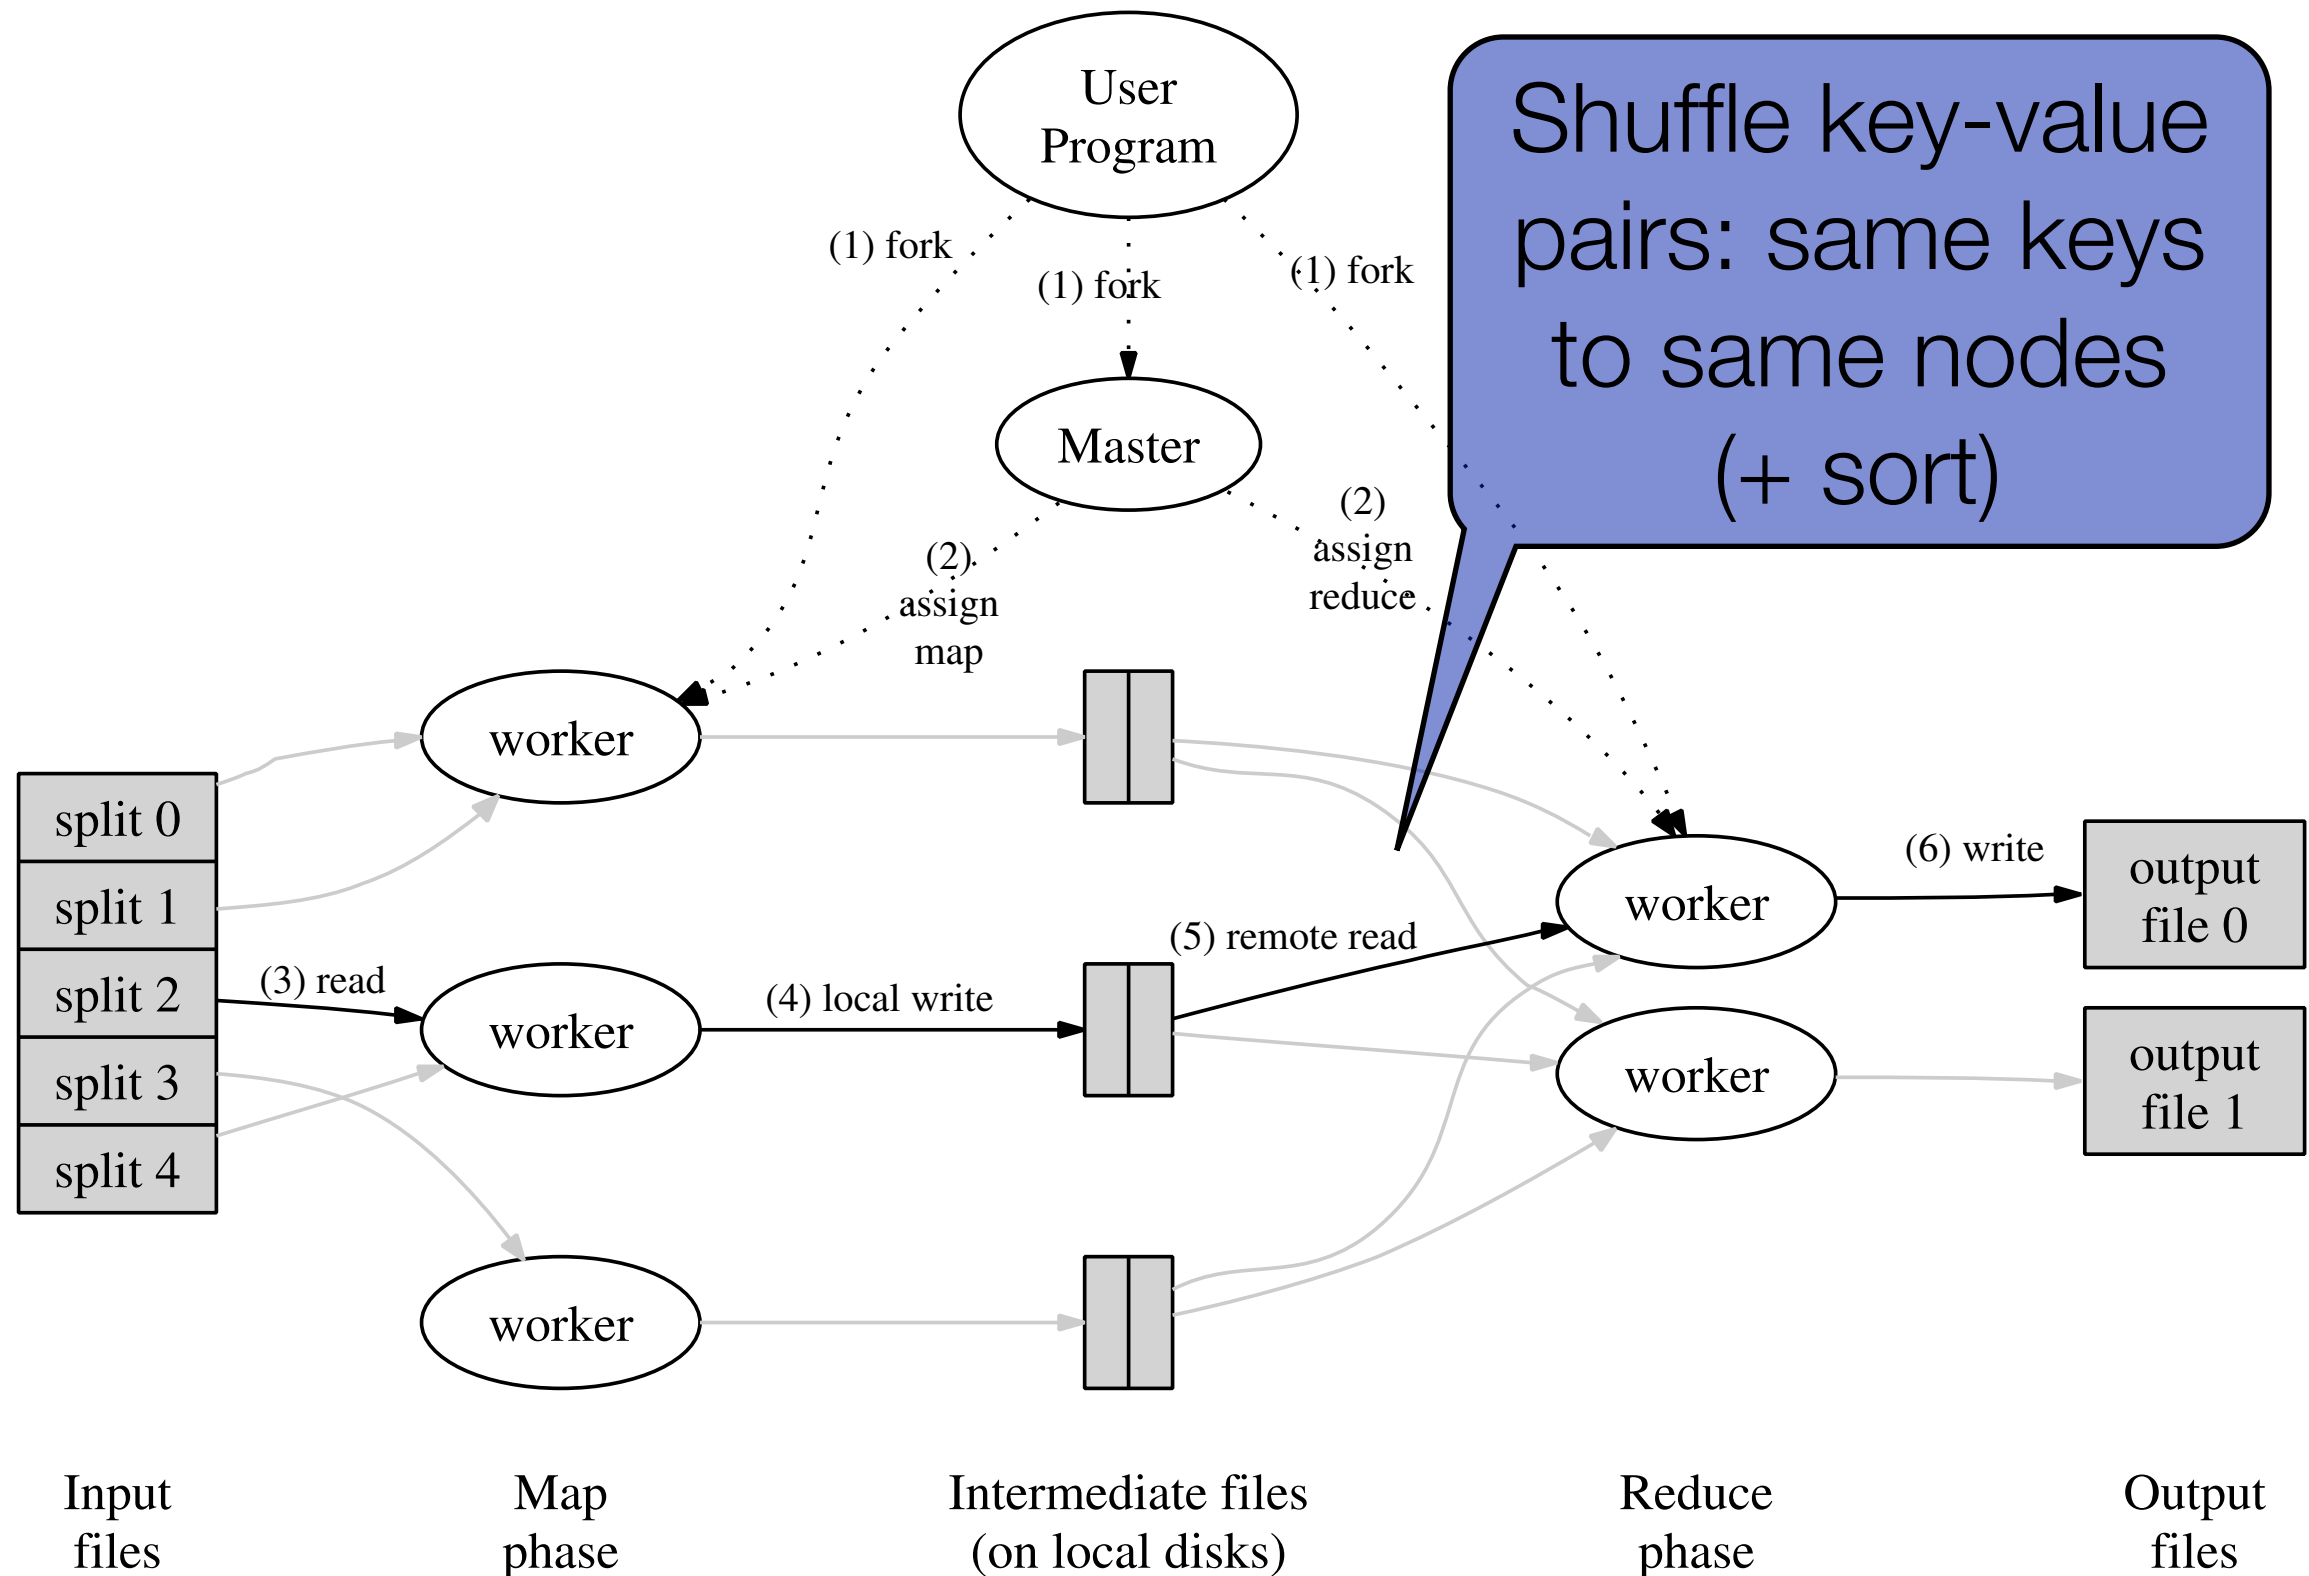
Map-reduce: distribute I/O over many nodes



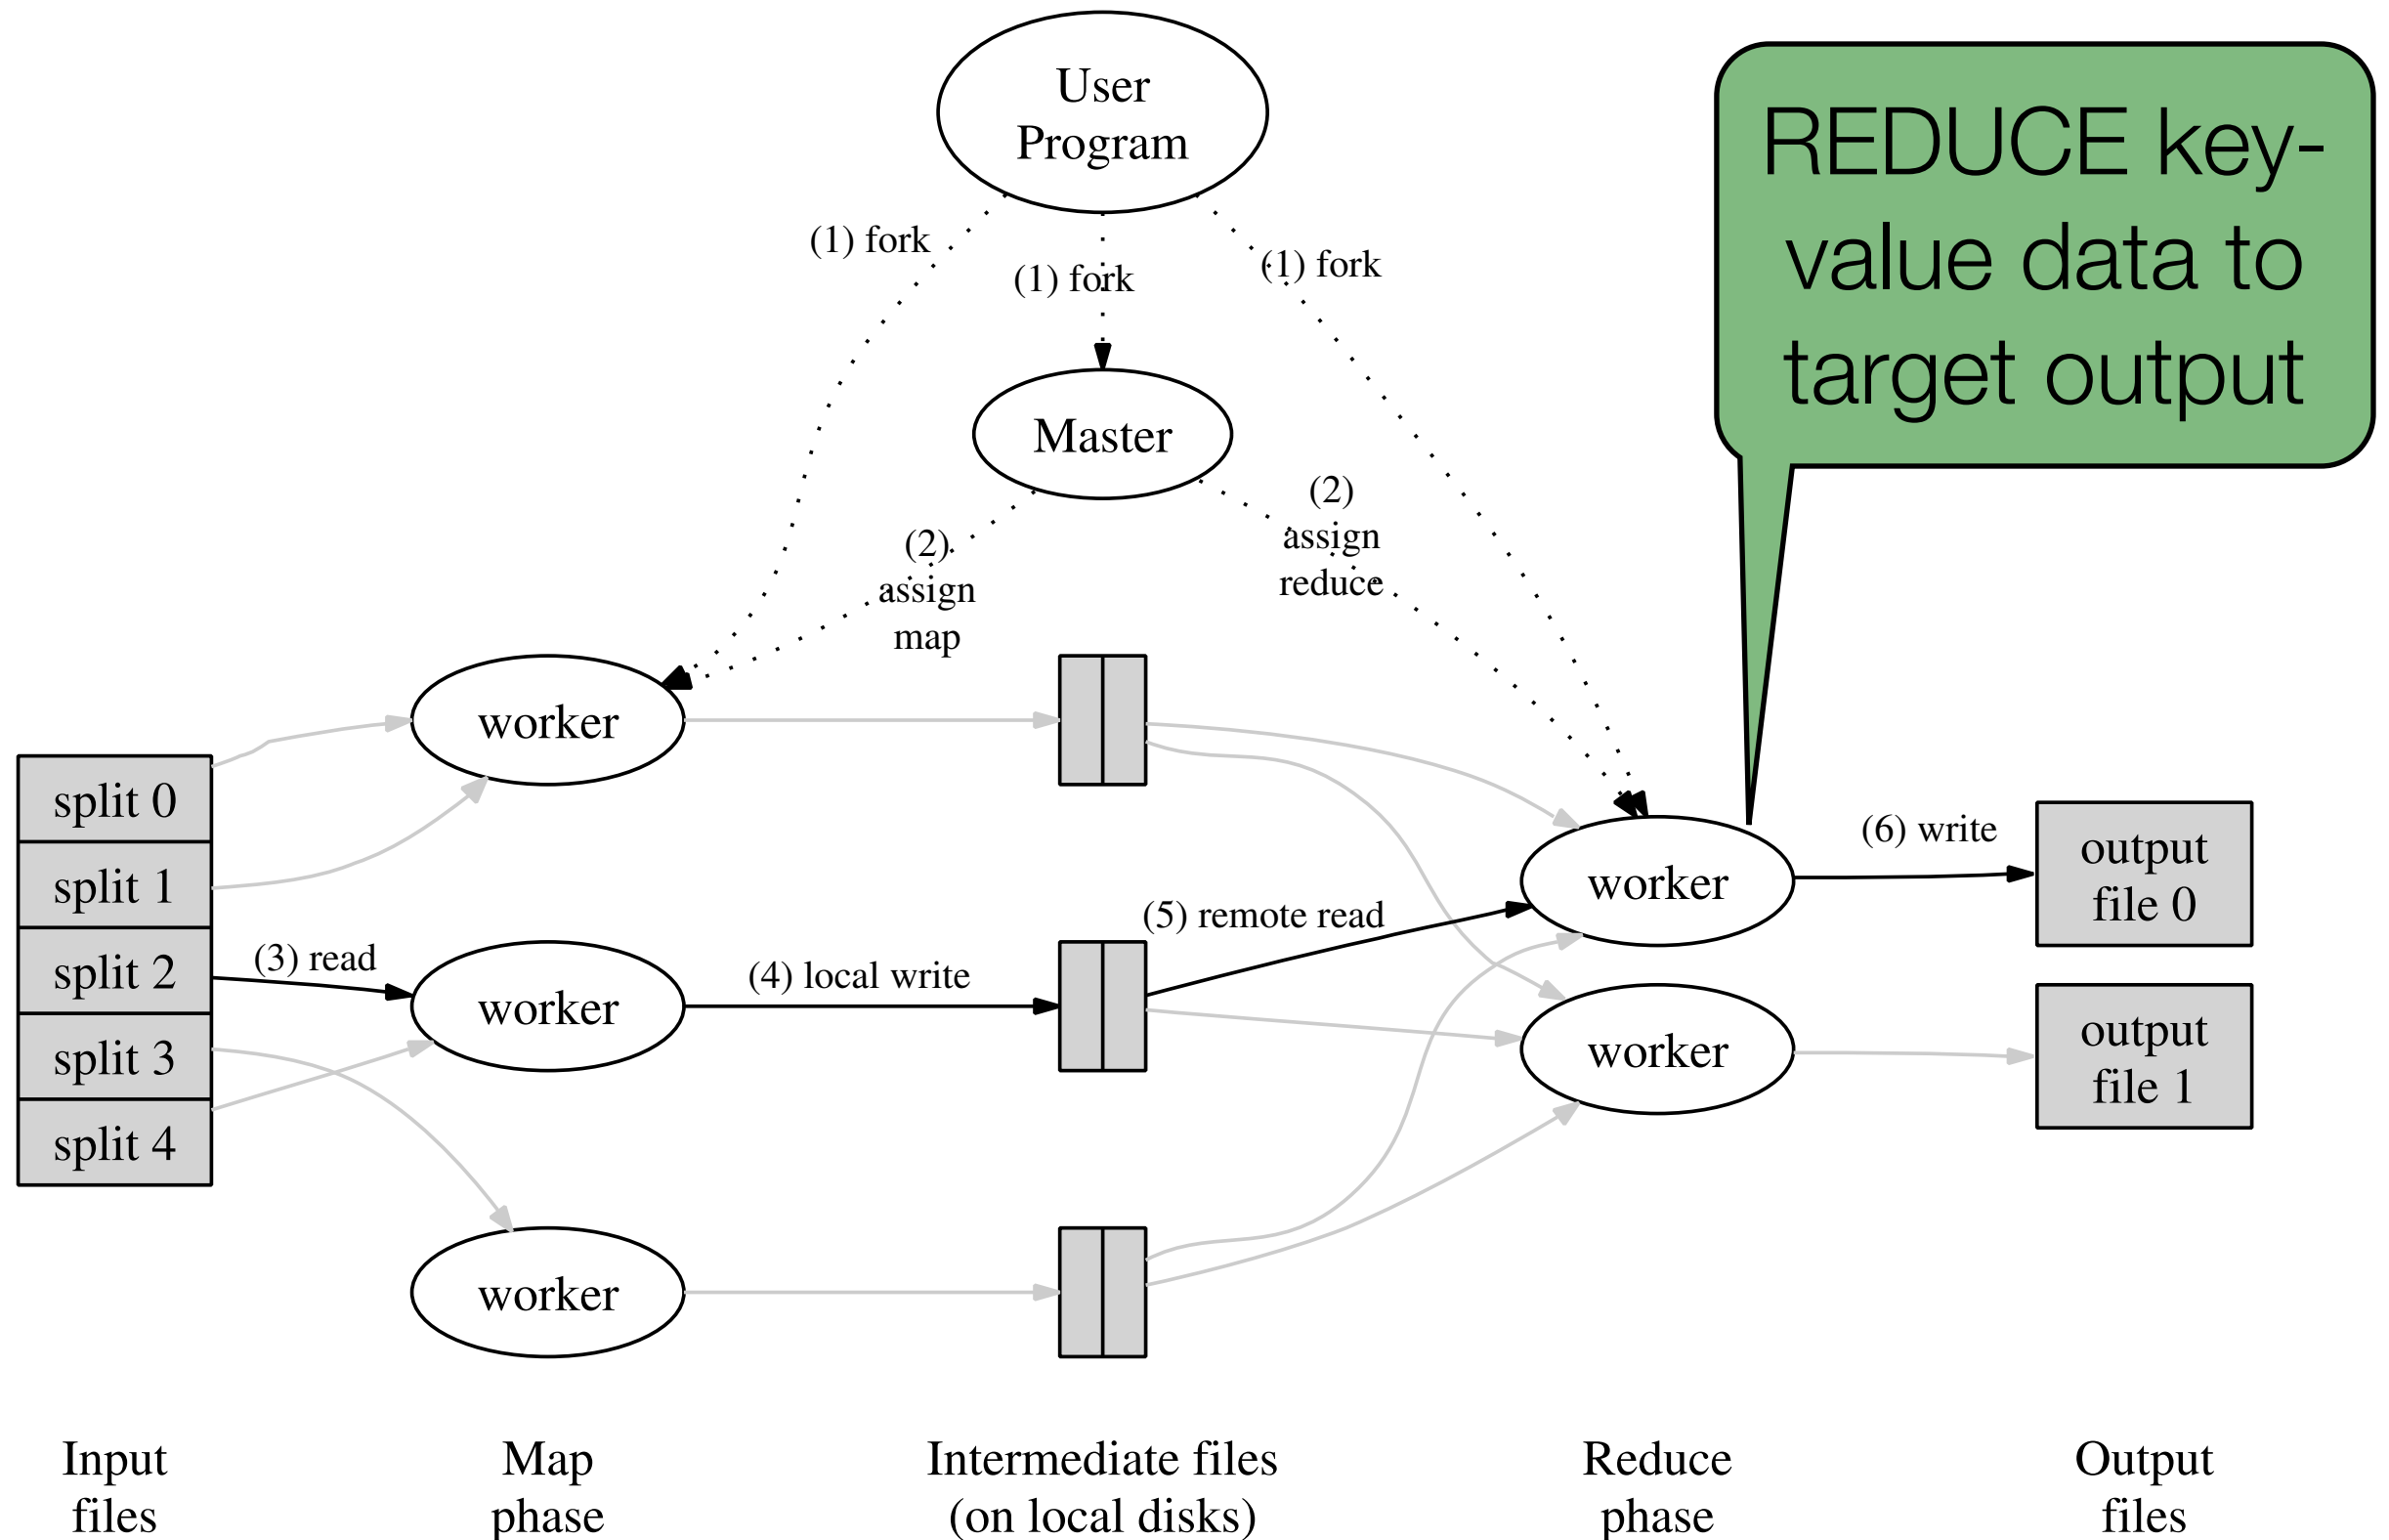
Map-reduce: distribute I/O over many nodes



Map-reduce: distribute I/O over many nodes




Map-reduce: distribute I/O over many nodes




Map-reduce: Intuition

One apple



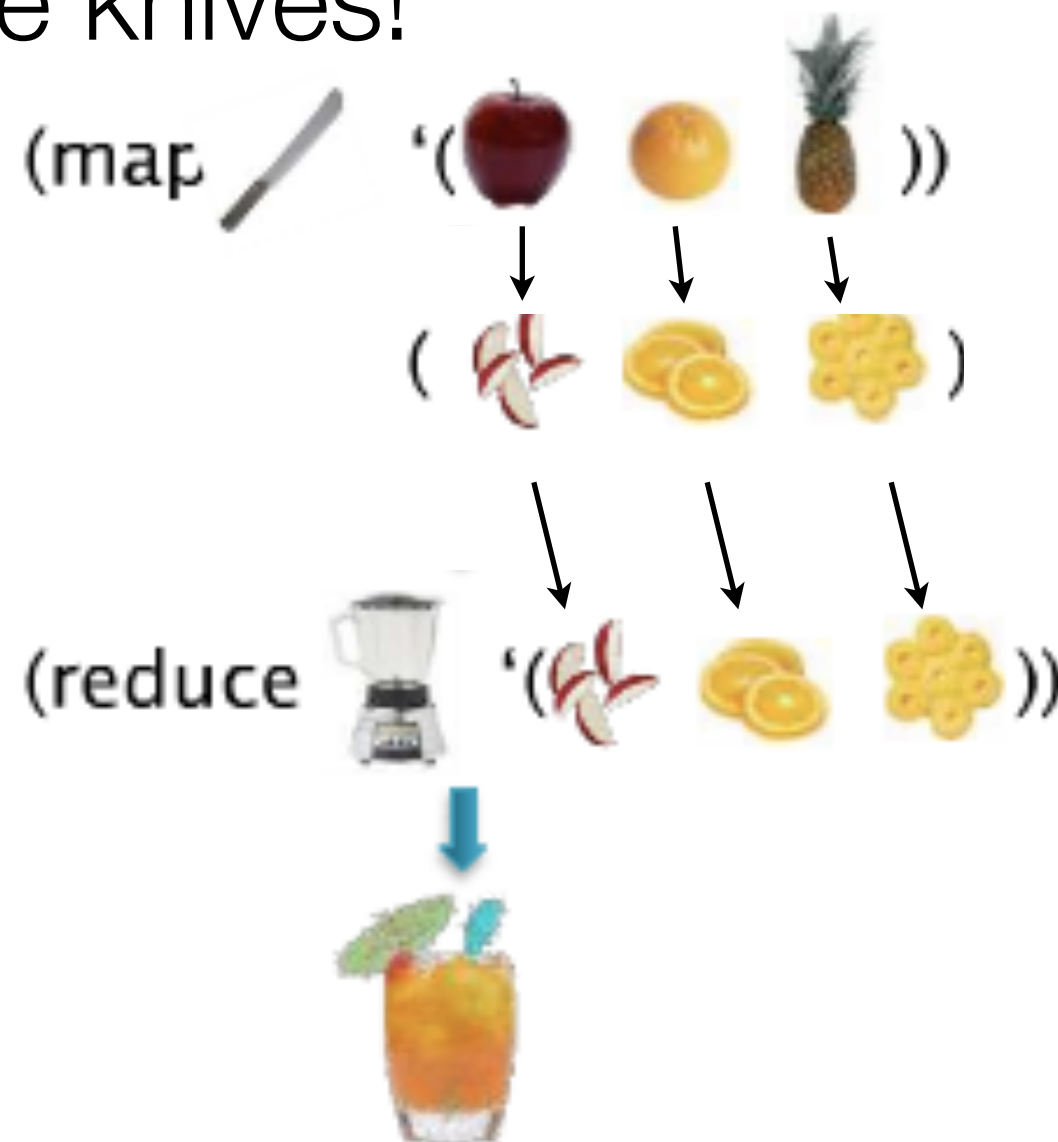
Map-reduce: Intuition

One apple



A fruit basket: more knives!


many mappers
(people with knives)



1 reducer
(person with blender)

Map-reduce: Intuition

A fruit container: more blenders too!




Map input:

$\langle a, \text{🍏} \rangle \quad \langle o, \text{🍊} \rangle \quad \langle p, \text{🍍} \rangle, \dots$

Map-reduce: Intuition

A fruit container: more blenders too!



Map input:


$\langle a, \text{apple} \rangle \langle o, \text{orange} \rangle \langle p, \text{pineapple} \rangle, \dots$

Map output:

$\langle a', \text{apple slices} \rangle \langle o', \text{orange slices} \rangle \langle p', \text{pineapple slices} \rangle, \dots$

Map-reduce: Intuition

A fruit container: more blenders too!



Map input:

$\langle a, \text{apple} \rangle \langle o, \text{orange} \rangle \langle p, \text{pineapple} \rangle, \dots$

Map output:


$\langle a', \text{apple slice} \rangle \langle o', \text{orange slice} \rangle \langle p', \text{pineapple slice} \rangle, \dots$

Reduce input:

$\langle a', \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice} \rangle$

Map-reduce: Intuition

A fruit container: more blenders too!



Map input:

$\langle a, \text{apple} \rangle \langle o, \text{orange} \rangle \langle p, \text{pineapple} \rangle, \dots$

Map output:

$\langle a', \text{apple slice} \rangle \langle o', \text{orange slice} \rangle \langle p', \text{pineapple slice} \rangle, \dots$

Reduce input:


$\langle a', \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice} \rangle$

Reduce output:



Map-reduce: Intuition

A fruit container: more blenders too!



Map input:

$\langle a, \text{apple} \rangle \langle o, \text{orange} \rangle \langle p, \text{pineapple} \rangle, \dots$

Map output:

$\langle a', \text{apple slice} \rangle \langle o', \text{orange slice} \rangle \langle p', \text{pineapple slice} \rangle, \dots$

Reduce input:


$\langle a', \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice} \rangle$

Reduce output:



Map-reduce: Intuition

A fruit container: more blenders too!



Map input:

$\langle a, \text{apple} \rangle \langle o, \text{orange} \rangle \langle p, \text{pineapple} \rangle, \dots$

Map output:

$\langle a', \text{apple slice} \rangle \langle o', \text{orange slice} \rangle \langle p', \text{pineapple slice} \rangle, \dots$

Reduce input:


$\langle a', \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice} \rangle$

Reduce output:



Map-reduce: Intuition

A fruit container: more blenders too!



Map input:

$\langle a, \text{apple} \rangle \langle o, \text{orange} \rangle \langle p, \text{pineapple} \rangle, \dots$

Map output:

$\langle a', \text{apple slice} \rangle \langle o', \text{orange slice} \rangle \langle p', \text{pineapple slice} \rangle, \dots$

Reduce input:


$\langle a', \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice} \rangle$

Reduce output:



Map-reduce: Intuition

A fruit container: more blenders too!



Map input:

$\langle a, \text{apple} \rangle \langle o, \text{orange} \rangle \langle p, \text{pineapple} \rangle, \dots$

Map output:

$\langle a', \text{apple slice} \rangle \langle o', \text{orange slice} \rangle \langle p', \text{pineapple slice} \rangle, \dots$

Reduce input:

$\langle a', \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice} \rangle$


Reduce output:



Map-reduce: Intuition

an academic analogy:
papers, reviewers, editors

A fruit container: more blenders too!



Map input:

$\langle a, \text{apple} \rangle \langle o, \text{orange} \rangle \langle p, \text{pineapple} \rangle, \dots$

Map output:

$\langle a', \text{apple slice} \rangle \langle o', \text{orange slice} \rangle \langle p', \text{pineapple slice} \rangle, \dots$

Reduce input:

$\langle a', \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice}, \text{apple slice} \rangle$

Reduce output:



Example: inverted indexing (e.g. webpages)

Example: inverted indexing (e.g. webpages)

- Input: documents/webpages
 - Doc 1: "Why did the chicken cross the road?"
 - Doc 2: "The chicken and egg problem"
 - Doc 3: "Kentucky Fried Chicken"
- Output
 - Index of words:
 - The word "the" occurs twice in Doc 1 (positions 3 and 6), once in Doc 2 (position 1)

Example: inverted indexing (e.g. webpages)

Map phase (3 parallel tasks)

- $\text{map}_1 \Rightarrow ("why", (\text{doc}_1, 1)), ("did", (\text{doc}_1, 2)), ("the", (\text{doc}_1, 3)),$
 $("chicken", (\text{doc}_1, 4)), ("cross", (\text{doc}_1, 5)), ("the", (\text{doc}_1, 6)),$
 $("road", (\text{doc}_1, 7))$
- $\text{map}_2 \Rightarrow ("the", (\text{doc}_2, 1)), ("chicken", (\text{doc}_2, 2)), ("and", (\text{doc}_2, 3)),$
 $("egg", (\text{doc}_2, 4)), ("problem", (\text{doc}_2, 5))$
- $\text{map}_3 \Rightarrow ("kentucky", (\text{doc}_3, 1)), ("fried", (\text{doc}_3, 2)), ("chicken", (\text{doc}_3, 3))$

Example: inverted indexing (e.g. webpages)

Map phase (3 parallel tasks)

- $\text{map}_1 \Rightarrow (\text{"why"}, (\text{doc}_1, 1)), (\text{"did"}, (\text{doc}_1, 2)), (\text{"the"}, (\text{doc}_1, 3)), (\text{"chicken"}, (\text{doc}_1, 4)), (\text{"cross"}, (\text{doc}_1, 5)), (\text{"the"}, (\text{doc}_1, 6)), (\text{"road"}, (\text{doc}_1, 7))$
- $\text{map}_2 \Rightarrow (\text{"the"}, (\text{doc}_2, 1)), (\text{"chicken"}, (\text{doc}_2, 2)), (\text{"and"}, (\text{doc}_2, 3)), (\text{"egg"}, (\text{doc}_2, 4)), (\text{"problem"}, (\text{doc}_2, 5))$
- $\text{map}_3 \Rightarrow (\text{"kentucky"}, (\text{doc}_3, 1)), (\text{"fried"}, (\text{doc}_3, 2)), (\text{"chicken"}, (\text{doc}_3, 3))$

Intermediate shuffle & sort phase

- $(\text{"why"}, \langle (\text{doc}_1, 1) \rangle)$,
- $(\text{"did"}, \langle (\text{doc}_1, 2) \rangle)$,
- $(\text{"the"}, \langle (\text{doc}_1, 3), (\text{doc}_1, 6), (\text{doc}_2, 1) \rangle)$
- $(\text{"chicken"}, \langle (\text{doc}_1, 4), (\text{doc}_2, 2), (\text{doc}_3, 3) \rangle)$
- $(\text{"cross"}, \langle (\text{doc}_1, 5) \rangle)$
- $(\text{"road"}, \langle (\text{doc}_1, 7) \rangle)$
- $(\text{"and"}, \langle (\text{doc}_2, 3) \rangle)$
- $(\text{"egg"}, \langle (\text{doc}_2, 4) \rangle)$
- $(\text{"problem"}, \langle (\text{doc}_2, 5) \rangle)$

Example: inverted indexing (e.g. webpages)

Intermediate shuffle & sort phase

- (“why”, <(doc₁,1)>),
- (“did”, <(doc₁,2)>),
- (“the”, <(doc₁,3), (doc₁,6), (doc₂,1)>)
- (“chicken”, <(doc₁,4), (doc₂,2), (doc₃,3)>)
- (“cross”, <(doc₁,5)>)
- (“road”, <(doc₁,7)>)
- (“and”, <(doc₂,3)>)
- (“egg”, <(doc₂,4)>)
- (“problem”, <(doc₂,5)>)
- (“kentucky”, <(doc₃,1)>)
- (“fried”, <(doc₃,2)>)

Example: inverted indexing (e.g. webpages)

Intermediate shuffle & sort phase

- (“why”, <(doc₁,1)>),
- (“did”, <(doc₁,2)>),
- (“the”, <(doc₁,3), (doc₁,6), (doc₂,1)>)
- (“chicken”, <(doc₁,4), (doc₂,2), (doc₃,3)>)
- (“cross”, <(doc₁,5)>)
- (“road”, <(doc₁,7)>)
- (“and”, <(doc₂,3)>)
- (“egg”, <(doc₂,4)>)
- (“problem”, <(doc₂,5)>)
- (“kentucky”, <(doc₃,1)>)
- (“fried”, <(doc₃,2)>)

Reduce phase (11 parallel tasks)

- (“why”, <(doc₁,<1>)>),
- (“did”, <(doc₁,<2>)>),
- (“the”, <(doc₁,<3,6>), (doc₂,<1>)>)
- (“chicken”, <(doc₁,<4>), (doc₂,<2>), (doc₃,<3>)>)
- (“cross”, <(doc₁,<5>)>)
- (“road”, <(doc₁,<7>)>)
- (“and”, <(doc₂,<3>)>)
- (“egg”, <(doc₂,<4>)>)
- (“problem”, <(doc₂,<5>)>)
- (“kentucky”, <(doc₃,<1>)>)
- (“fried”, <(doc₃,<2>)>)

Example: inverted indexing (e.g. webpages)

Intermediate shuffle & sort phase

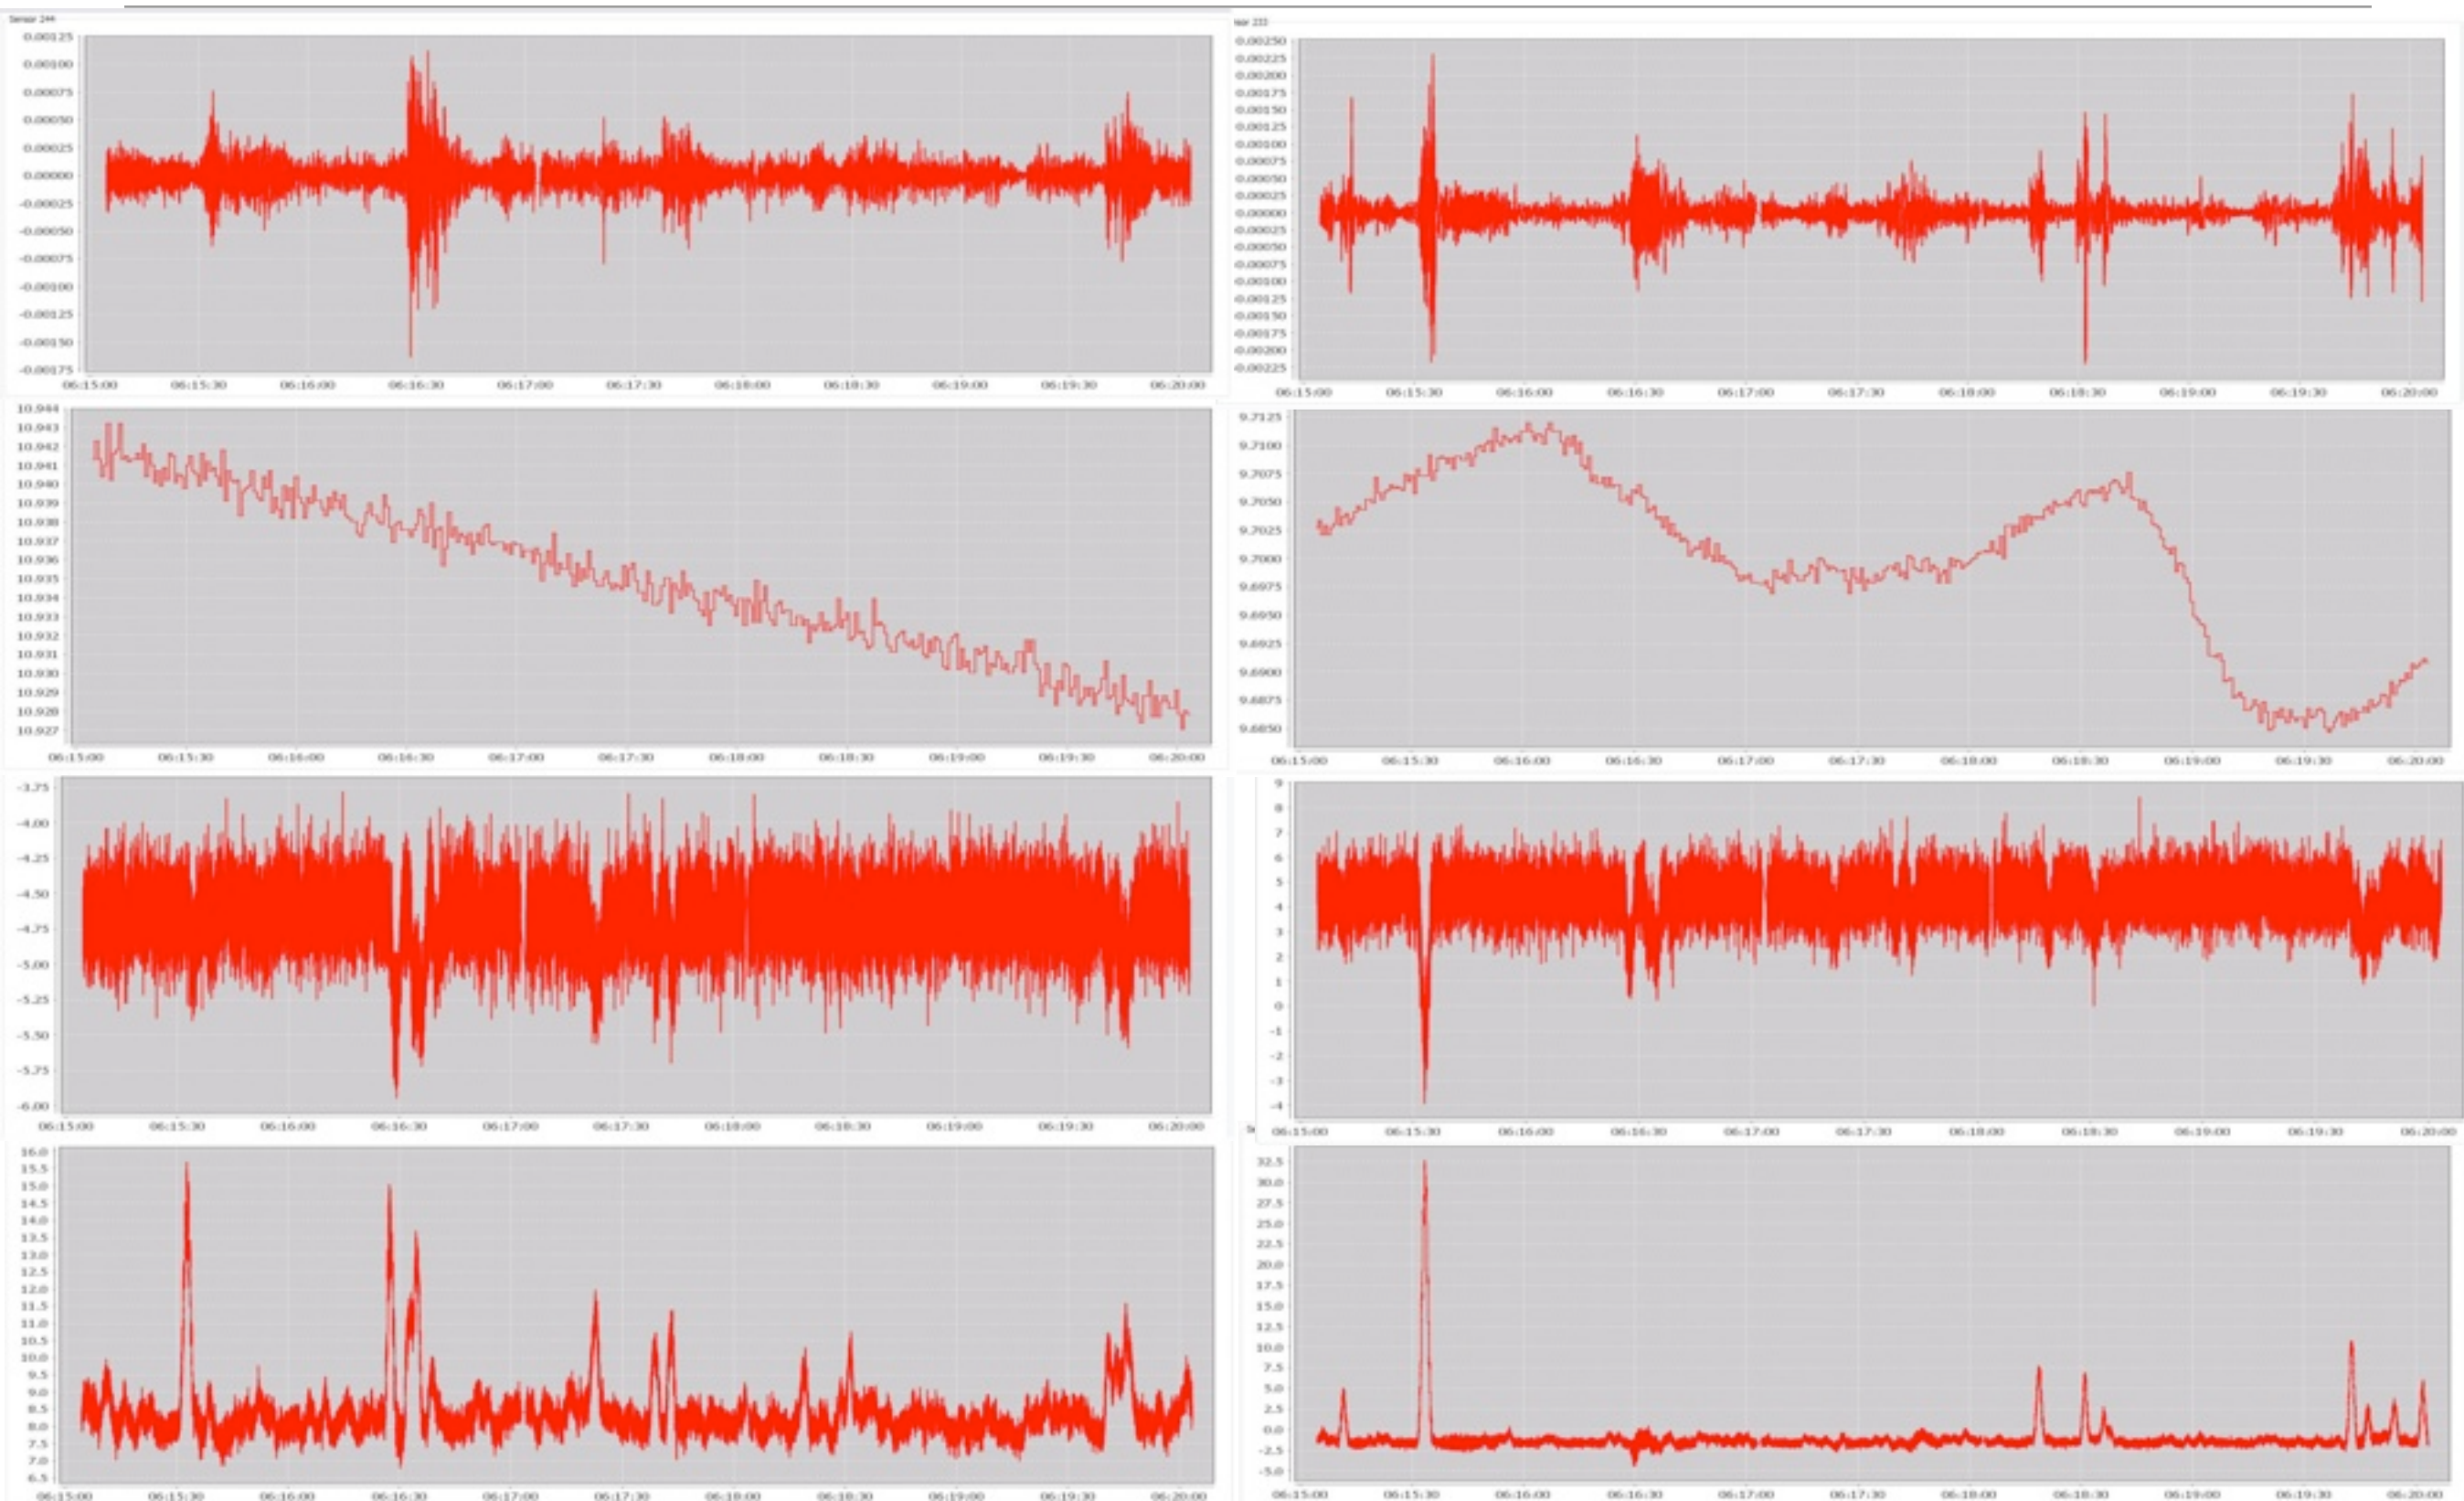
- (“why”, <(doc₁,1)>),
- (“did”, <(doc₁,2)>),
- (“the”, <(doc₁,3), (doc₁,6), (doc₂,1)>)
- (“chicken”, <(doc₁,4), (doc₂,2), (doc₃,3)>)
- (“cross”, <(doc₁,5)>)
- (“road”, <(doc₁,7)>)
- (“and”, <(doc₂,3)>)
- (“egg”, <(doc₂,4)>)
- (“problem”, <(doc₂,5)>)
- (“kentucky”, <(doc₃,1)>)
- (“fried”, <(doc₃,2)>)

Reduce phase (11 parallel tasks)

- (“why”, <(doc₁,<1>)>),
- (“did”, <(doc₁,<2>)>),
- (“the”, <(doc₁,<3,6>), (doc₂,<1>)>)
- (“chicken”, <(doc₁,<4>), (doc₂,<2>), (doc₃,<3>)>)
- (“cross”, <(doc₁,<5>)>)
- (“road”, <(doc₁,<7>)>)
- (“and”, <(doc₂,<3>)>)
- (“egg”, <(doc₂,<4>)>)
- (“problem”, <(doc₂,<5>)>)
- (“kentucky”, <(doc₃,<1>)>)
- (“fried”, <(doc₃,<2>)>)

Simply write mapper and reducer method, often few lines of code

Map-reduce on time series data



A simple operation: aggregation

2008-10-24 06:15:04.559, -6.293695, -1.1263204, 2.985364, 43449.957, 2.3577218, 38271.21
2008-10-24 06:15:04.570, -6.16952, -1.3805857, 2.6128333, 43449.957, 2.4848552, 37399.26
2008-10-24 06:15:04.580, -5.711255, -0.8897944, 3.139107, 43449.957, 2.1744132, 38281.0

2008-10-24 06:15:04.559 min -524.0103
2008-10-24 06:15:04.559 max 38271.21
2008-10-24 06:15:04.559 avg 447.37795925930226
2008-10-24 06:15:04.570 min -522.7882
2008-10-24 06:15:04.570 max 37399.26
2008-10-24 06:15:04.570 avg 437.1266600847675

A simple operation: aggregation

- Input: Table with sensors in columns and timestamps in rows

2008-10-24 06:15:04.559, -6.293695, -1.1263204, 2.985364, 43449.957, 2.3577218, 38271.21
2008-10-24 06:15:04.570, -6.16952, -1.3805857, 2.6128333, 43449.957, 2.4848552, 37399.26
2008-10-24 06:15:04.580, -5.711255, -0.8897944, 3.139107, 43449.957, 2.1744132, 38281.0

- Desired output: Aggregated measures per timestamp

2008-10-24 06:15:04.559 min -524.0103
2008-10-24 06:15:04.559 max 38271.21
2008-10-24 06:15:04.559 avg 447.37795925930226
2008-10-24 06:15:04.570 min -522.7882
2008-10-24 06:15:04.570 max 37399.26
2008-10-24 06:15:04.570 avg 437.1266600847675

Map-reduce on time series data

```
public void map(LongWritable key, Text value, Context context) {
    String values[] = value.toString().split("\\t");
    Text time = new Text(values[0]);
    for(int i = 1; i <= nrStressSensors; i++)
        context.write(time, new Text(values[i]));
}

public void reduce(Text key, Iterable<Text> values, Context context) {
    //init; sum, min, max, count = 0
    Double d;
    for (Text v : values) {
        d = Double.valueOf(v.toString());
        sum += d;
        min = Math.min(min, d);
        max = Math.max(max, d);
        count++;
    }
    context.write(new Text(key+" min"), new Text(Double.toString((min))));
    context.write(new Text(key+" max"), new Text(Double.toString((max))));
    context.write(new Text(key+" avg"), new Text(Double.toString((sum/count))));
}
```


Map-reduce on time series data

```
public void map(LongWritable key, Text value, Context context) {  
    String values[] = value.toString().split("\\t");  
    Text time = new Text(values[0]);  
    for(int i = 1; i <= nrStressSensors; i++)  
        context.write(time, new Text(values[i]));  
}
```

MAP data to
<timestamp,value>

```
public void reduce(Text key, Iterable<Text> values, Context context) {  
    //init; sum, min, max, count = 0  
    Double d;  
    for (Text v : values) {  
        d = Double.valueOf(v.toString());  
        sum += d;  
        min = Math.min(min, d);  
        max = Math.max(max, d);  
        count++;  
    }  
    context.write(new Text(key+" min"), new Text(Double.toString((min))));  
    context.write(new Text(key+" max"), new Text(Double.toString((max))));  
    context.write(new Text(key+" avg"), new Text(Double.toString((sum/count))));  
}
```

Map-reduce on time series data

```
public void map(LongWritable key, Text value, Context context) {  
    String values[] = value.toString().split("\\t");  
    Text time = new Text(values[0]);  
    for(int i = 1; i <= nrStressSensors; i++)  
        context.write(time, new Text(values[i]));  
}
```

MAP data to
<timestamp,value>

```
public void reduce(Text key, Iterable<Text> values, Context context) {  
    //init; sum, min, max, count = 0  
    Double d;  
    for (Text v : values) {  
        d = Double.valueOf(v.toString());  
        sum += d;  
        min = Math.min(min, d);  
        max = Math.max(max, d);  
        count++;  
    }  
    context.write(new Text(key+" min"), new Text(Double.toString((min))));  
    context.write(new Text(key+" max"), new Text(Double.toString((max))));  
    context.write(new Text(key+" avg"), new Text(Double.toString((sum/count))));  
}
```

Shuffle/sort
per timestamp

Map-reduce on time series data

```
public void map(LongWritable key, Text value, Context context) {  
    String values[] = value.toString().split("\\t");  
    Text time = new Text(values[0]);  
    for(int i = 1; i <= nrStressSensors; i++)  
        context.write(time, new Text(values[i]));  
}
```

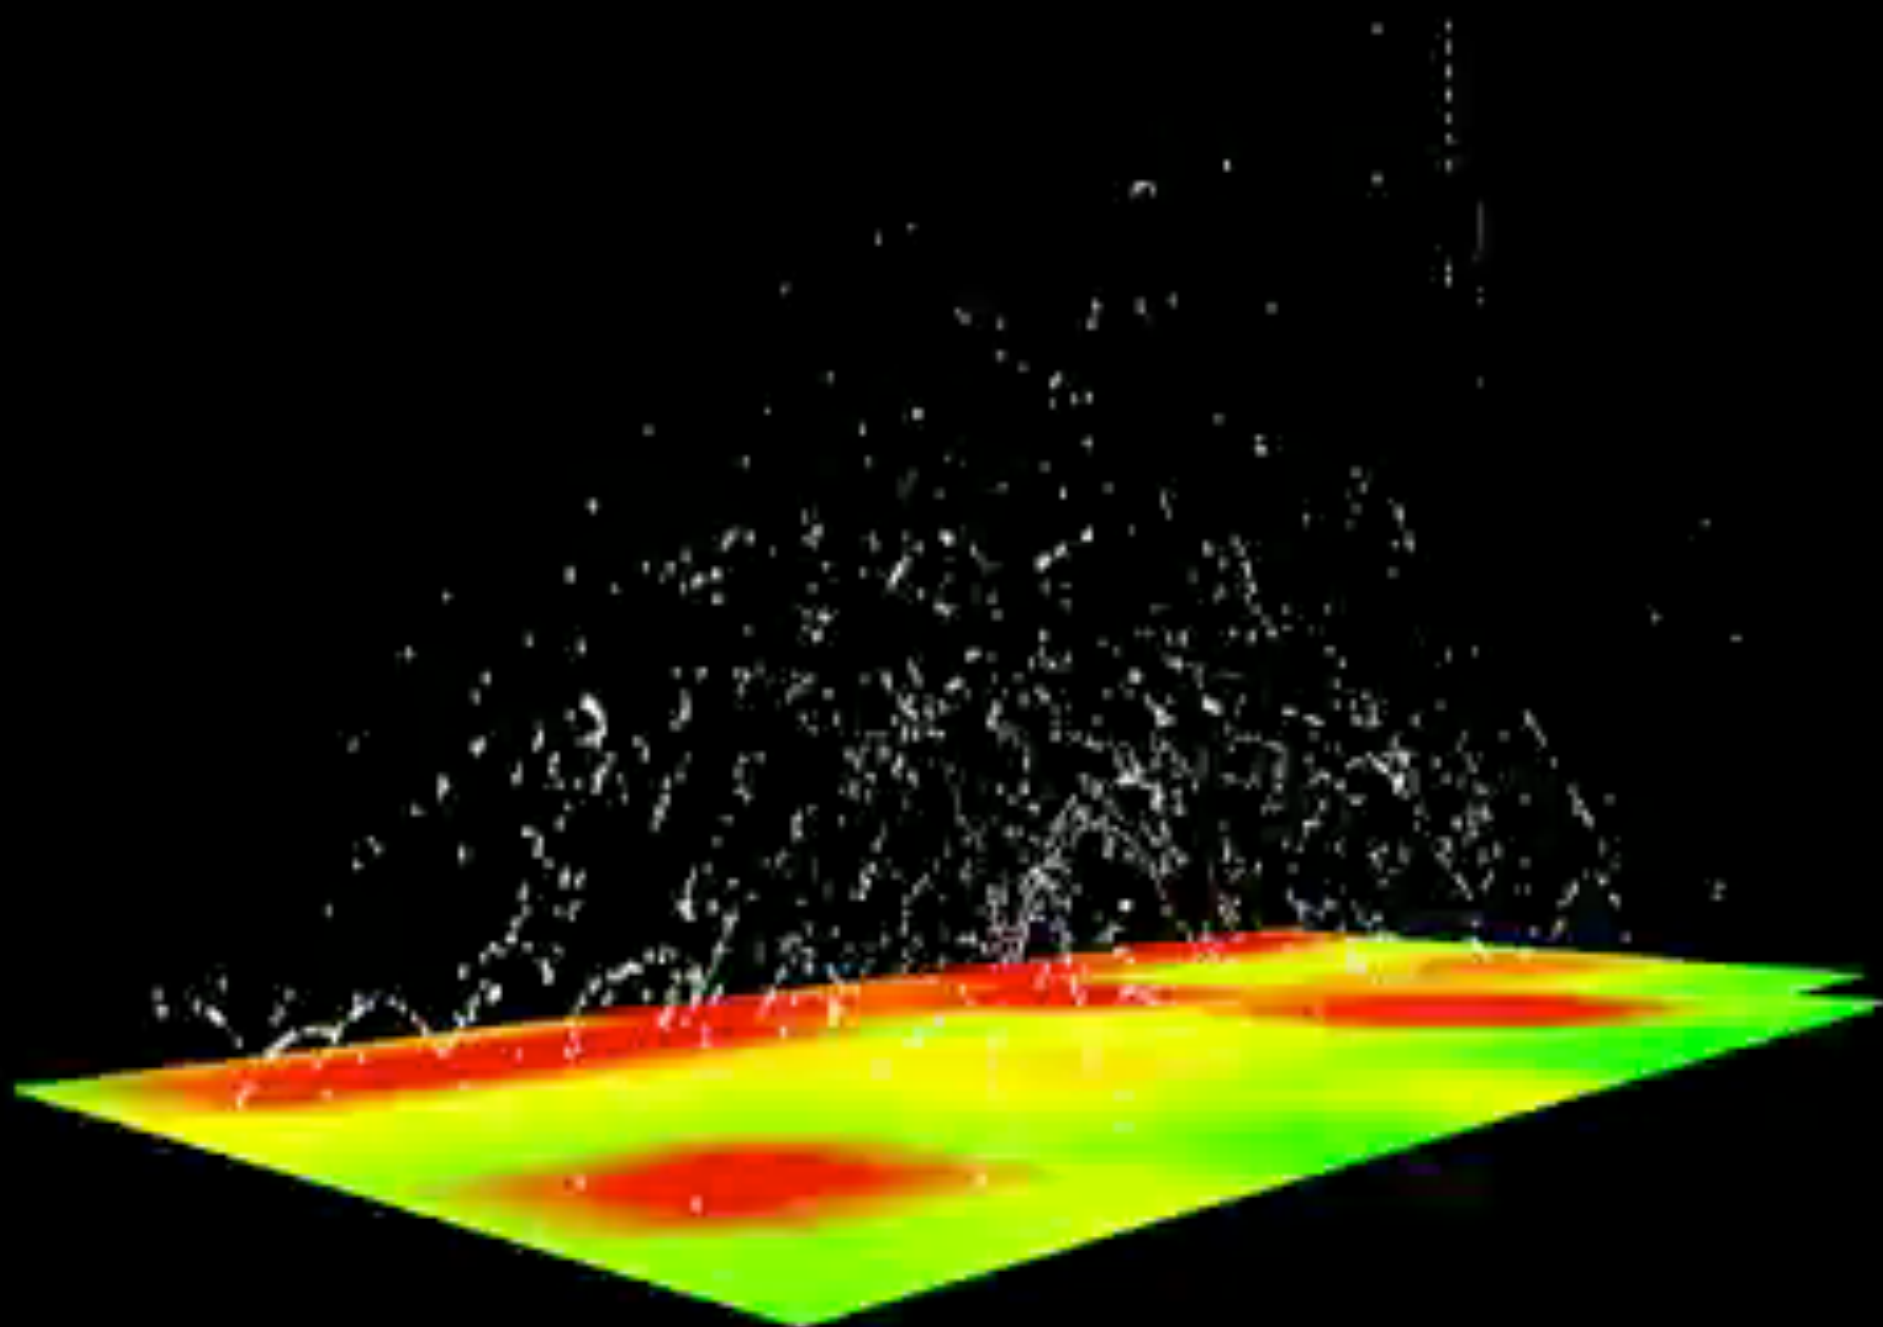
MAP data to
<timestamp,value>

```
public void reduce(Text key, Iterable<Text> values, Context context) {  
    //init; sum, min, max, count = 0  
    Double d;  
    for (Text v : values) {  
        d = Double.valueOf(v.toString());  
        sum += d;  
        min = Math.min(min, d);  
        max = Math.max(max, d);  
        count++;  
    }  
    context.write(new Text(key+" min"), new Text(Double.toString((min))));  
    context.write(new Text(key+" max"), new Text(Double.toString((max))));  
    context.write(new Text(key+" avg"), new Text(Double.toString((sum/count))));  
}
```

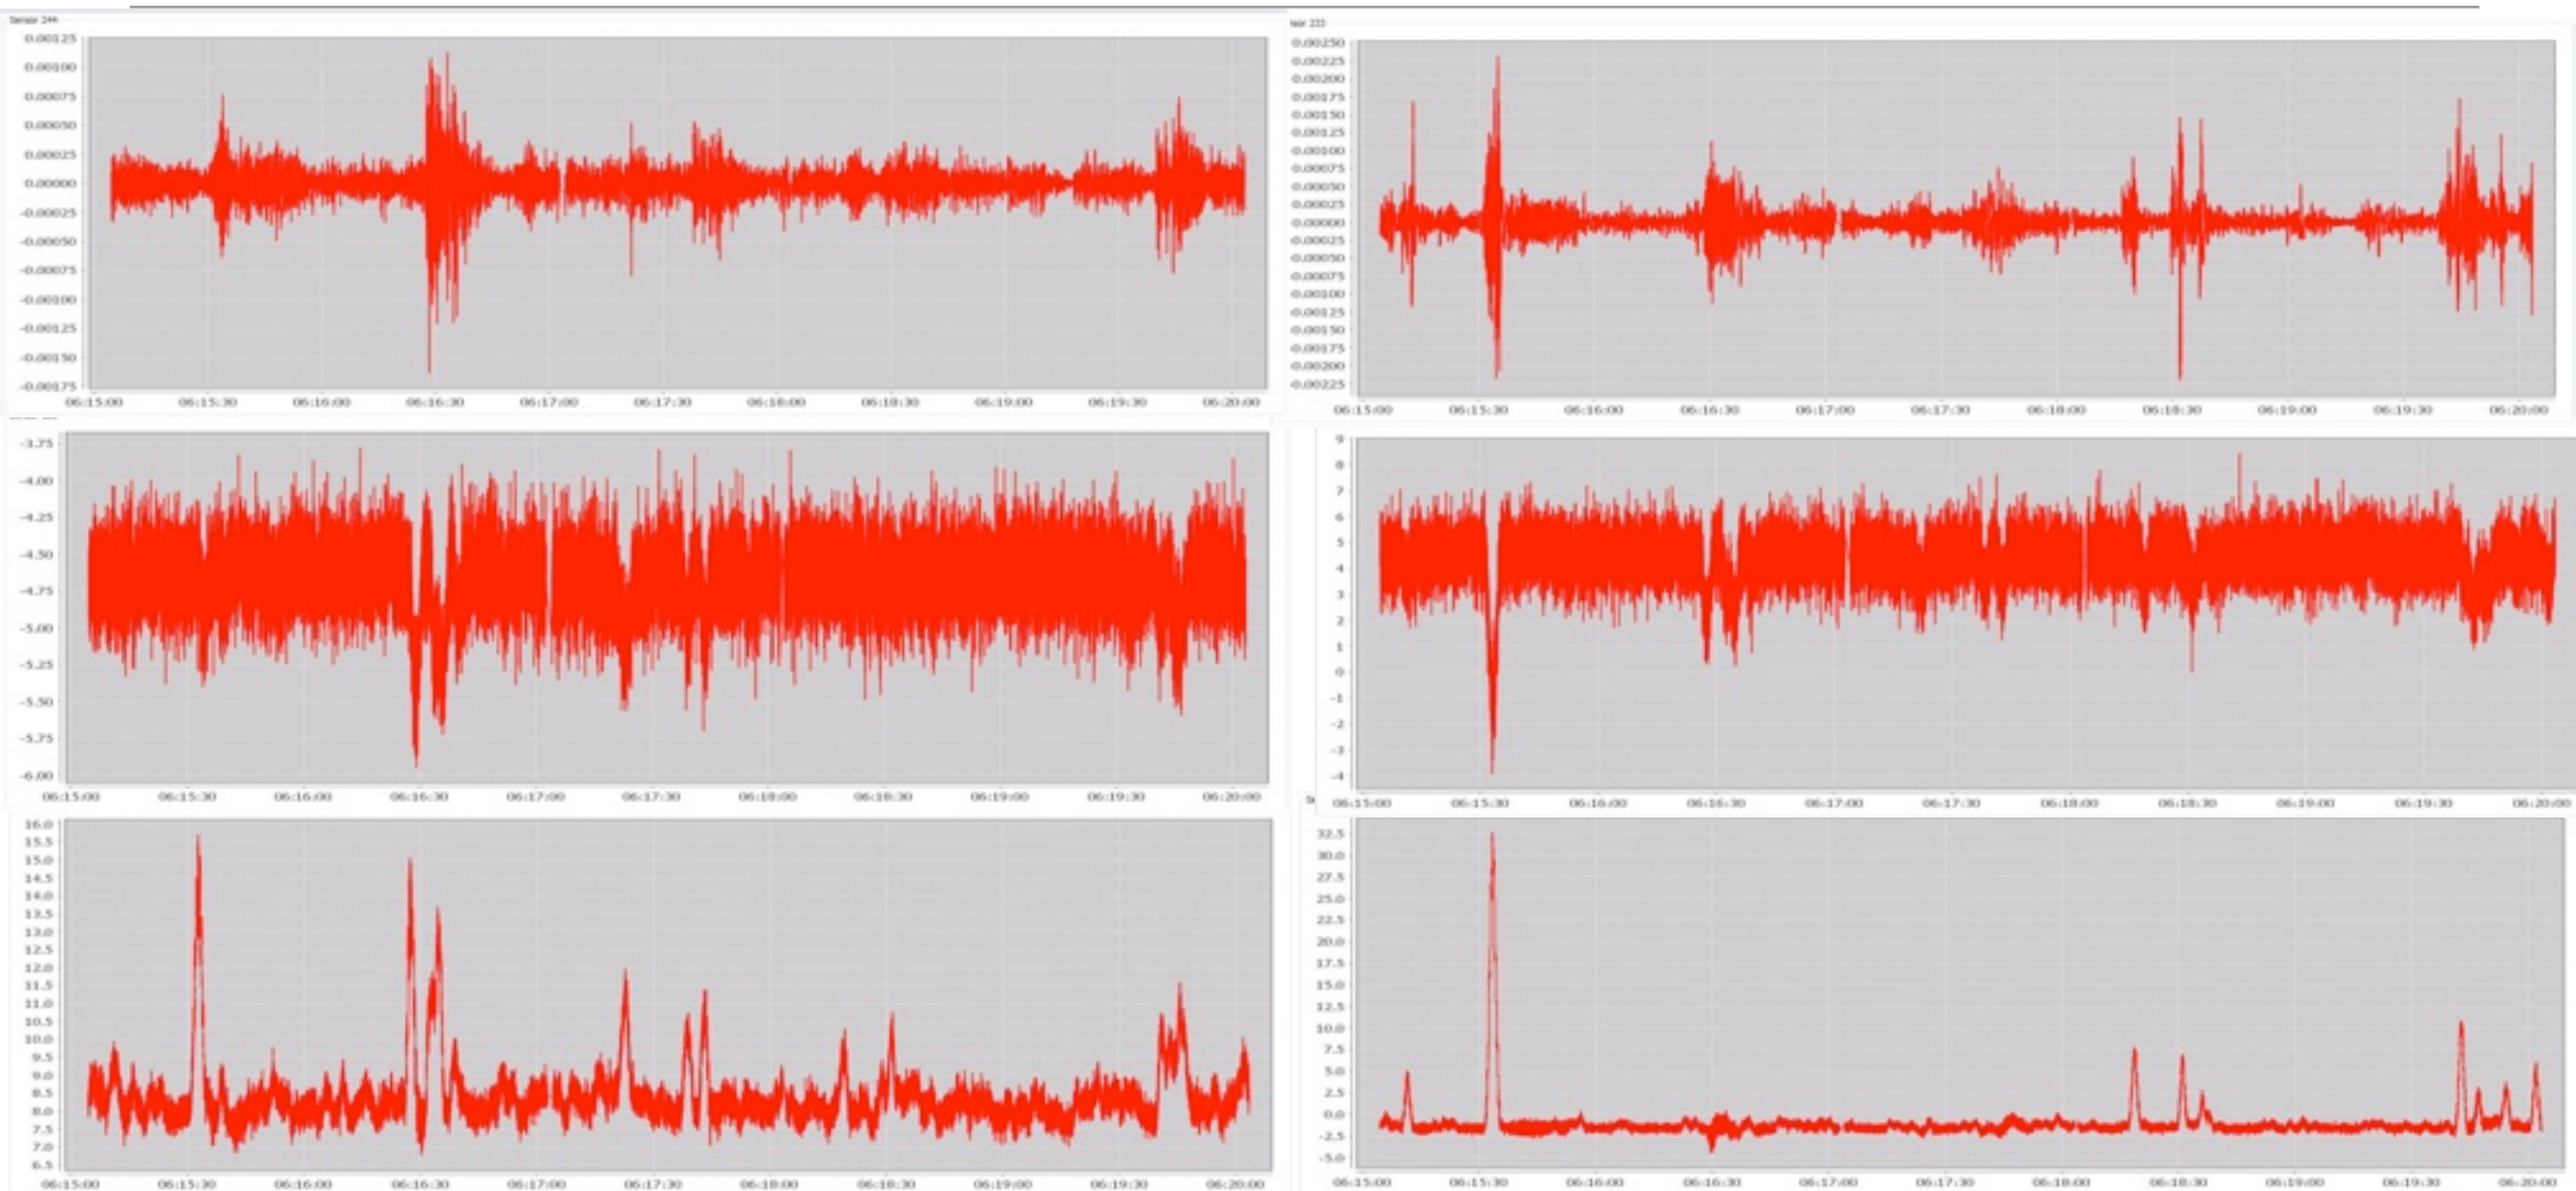
Shuffle/sort
per timestamp

REDUCE data per
timestamp to
aggregates

Shuffling




Map-reduce on time series data




Map-reduce on time series data


Min



Avg



Max

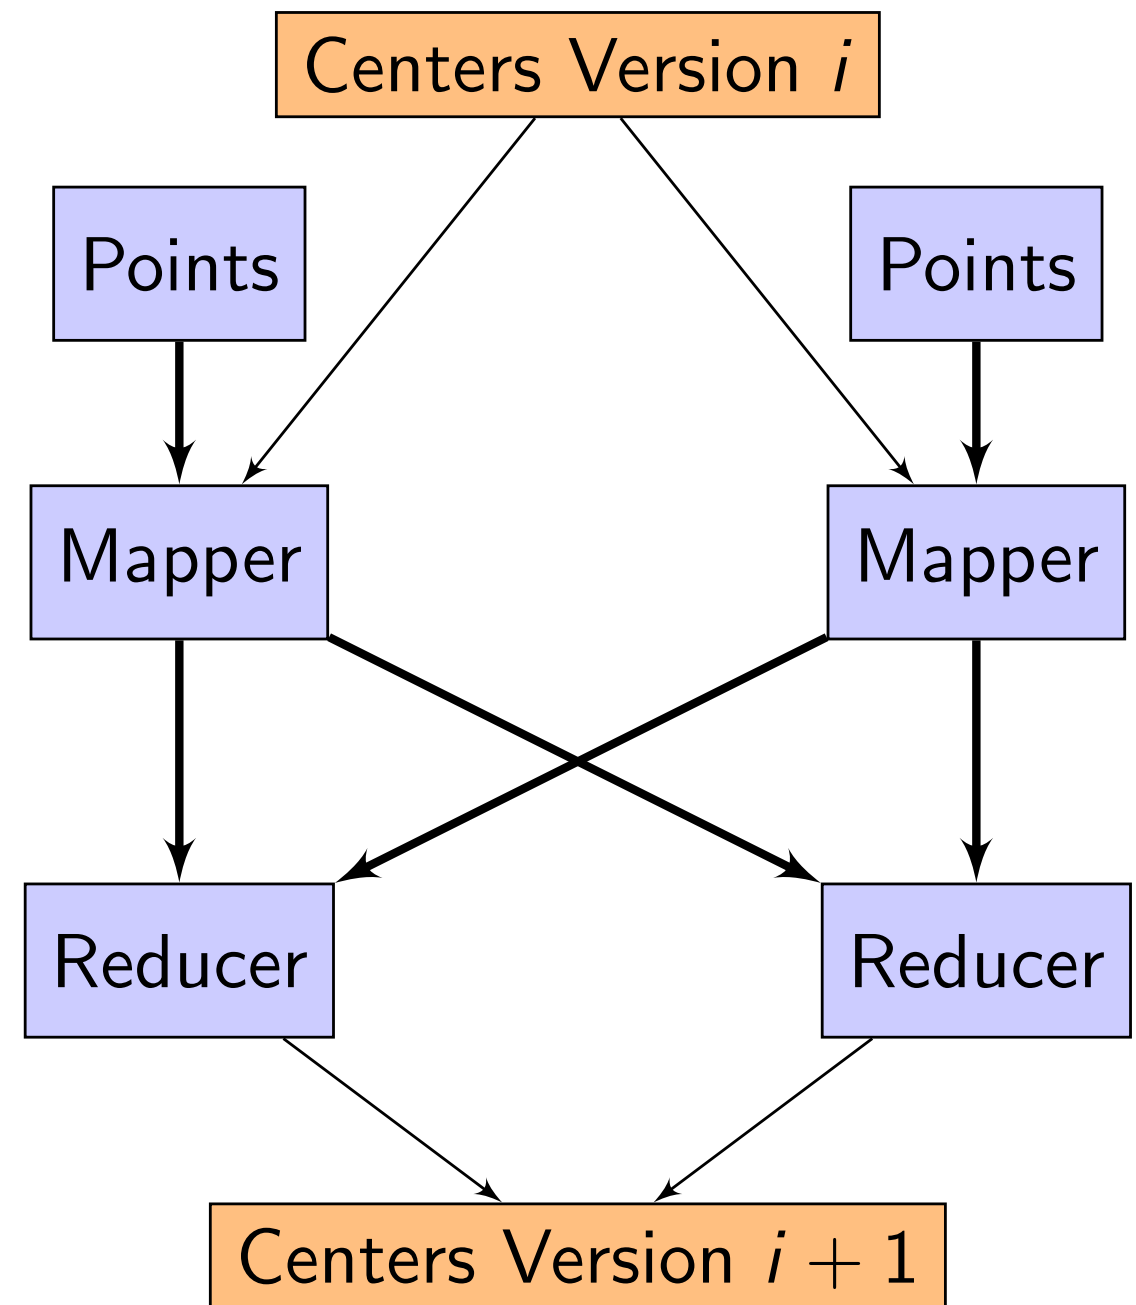


ML with map-reduce

- What part is I/O intensive (related to data points), and can be parallelized?
 - E.g. k-Means?

ML with map-reduce

- What part is I/O intensive (related to data points), and can be parallelized?
 - E.g. k-Means?
- Calculation of distance function!
 - Split data in chunks
 - Choose initial centers
 - MAP: calculate distances to all centers: $\langle \text{point}, [\text{distances}] \rangle$
 - REDUCE: calculate new centers
 - Repeat
- Others: SVM, NB, NN, LR,...
 - *Chu et al. Map-Reduce for ML on Multicore, NIPS '06*

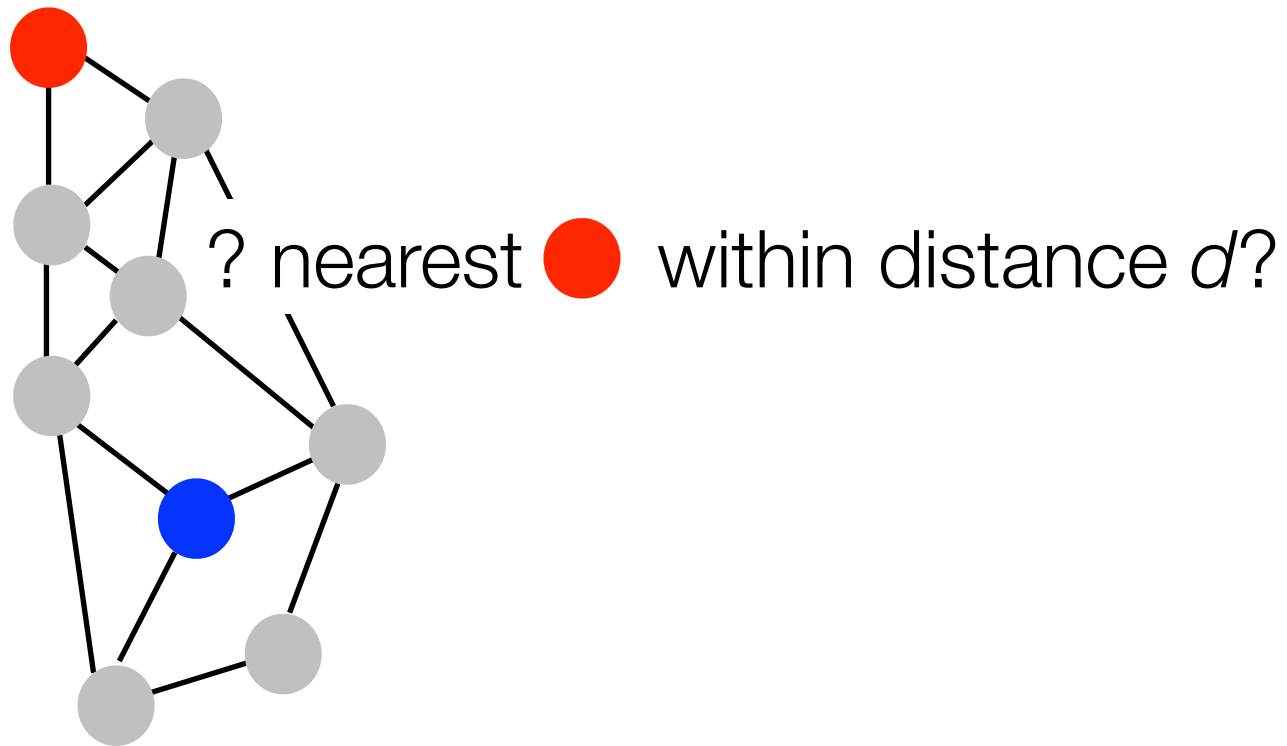


Map-reduce on graph data

- Find nearest feature on a graph

Input

graph
(node,label)



Map-reduce on graph data



- Find nearest feature on a graph

Input

graph
(node,label)

Map

\forall , search graph with radius d
 $< \text{}, \{ \text{}, distance \} >$



Map-reduce on graph data



- Find nearest feature on a graph

Input

graph
(node,label)

Map

\forall , search graph with radius d
 $< \text{}, \{ \text{}, distance \} >$



Map-reduce on graph data



- Find nearest feature on a graph

Input

graph
(node,label)

Map

\forall , search graph with radius d
 $< \text{}, \{ \text{}, distance \} >$



Map-reduce on graph data



- Find nearest feature on a graph

Input

graph
(node,label)

Map

\forall , search graph with radius d
< , { , $distance$ } >



Map-reduce on graph data



- Find nearest feature on a graph

Input

graph
(node,label)

Map

\forall , search graph with radius d
 $< \text{}, \{ \text{}, distance \} >$



Map-reduce on graph data



- Find nearest feature on a graph

Input

graph
(node,label)

Map

\forall , search graph with radius d
< , { , $distance$ } >






Map-reduce on graph data

- Find nearest feature on a graph

Input

graph
(node,label)


Map

\forall , search graph
< , { , *distance* } >

Shuffle/

Sort

by  id




Map-reduce on graph data

- Find nearest feature on a graph


Input

graph
(node,label)



Map

\forall red, search graph
 $< \text{gray}, \{\text{red}, \text{distance}\} >$



Shuffle/


Sort

by gray id



Reduce

$< \text{gray}, [\{\text{red}, \text{distance}\}, \{\text{blue}, \text{distance}\}] >$
 $\rightarrow \min()$




Map-reduce on graph data

- Find nearest feature on a graph


Input

graph
(node,label)



Map

\forall red, search graph
 $\langle \text{gray}, \{\text{red}, \text{distance}\} \rangle$




Shuffle/ Sort

by gray id




Reduce

$\langle \text{gray}, [\{\text{red}, \text{distance}\}, \{\text{blue}, \text{distance}\}] \rangle$
 $\rightarrow \min()$



Output

$\langle \text{gray}, \text{red} \rangle$
 $\langle \text{gray}, \text{blue} \rangle$
marked graph




Map-reduce on graph data

- Find nearest feature on a graph


Input

graph
(node,label)



Map

\forall red, search graph
 $\langle \text{gray}, \{\text{red}, \text{distance}\} \rangle$




Shuffle/ Sort

by gray id




Reduce

$\langle \text{gray}, [\{\text{red}, \text{distance}\}, \{\text{blue}, \text{distance}\}] \rangle$
 $\rightarrow \min()$



Output

$\langle \text{gray}, \text{red} \rangle$
 $\langle \text{gray}, \text{blue} \rangle$
marked graph



Be smart about mapping: load nearby nodes on same computing node: choose a good representation

The How?



Hadoop: The Definitive Guide
(Tom White, Yahoo!)

The How?



Hadoop: The Definitive Guide
(Tom White, Yahoo!)



Amazon Elastic Cloud (EC2)
Amazon Simple Storage Service (S3)
\$0.085/CPU hour, \$0.1/GBmonth

or... your university's HPC center
or... install your own cluster

The How?



Hadoop: The Definitive Guide
(Tom White, Yahoo!)



Amazon Elastic Cloud (EC2)
Amazon Simple Storage Service (S3)
\$0.085/CPU hour, \$0.1/GBmonth



Hadoop-based ML library

Classification: LR, NB, SVM*, NN*, RF
Clustering: kMeans, canopy, EM, spectral
Regression: LWLR*
Pattern mining: Top-k FPGrowth
Dim. Reduction, Coll. Filtering
(*under development)

or... your university's HPC center
or... install your own cluster


Part II: Grid computing (CPU bound)

- Disambiguation:
 - Supercomputer (shared memory)
 - CPU + memory bound, you're rich
 - Identical nodes
 - Cluster computing
 - Parallelizable over many nodes (MPI), you're rich/patient
 - Similar nodes
 - Grid computing
 - Parallelizable over few nodes or *embarrassingly parallel*
 - Very heterogenous nodes, but loads of `free' computing power

Grid computing



Grid computing



Grid computing

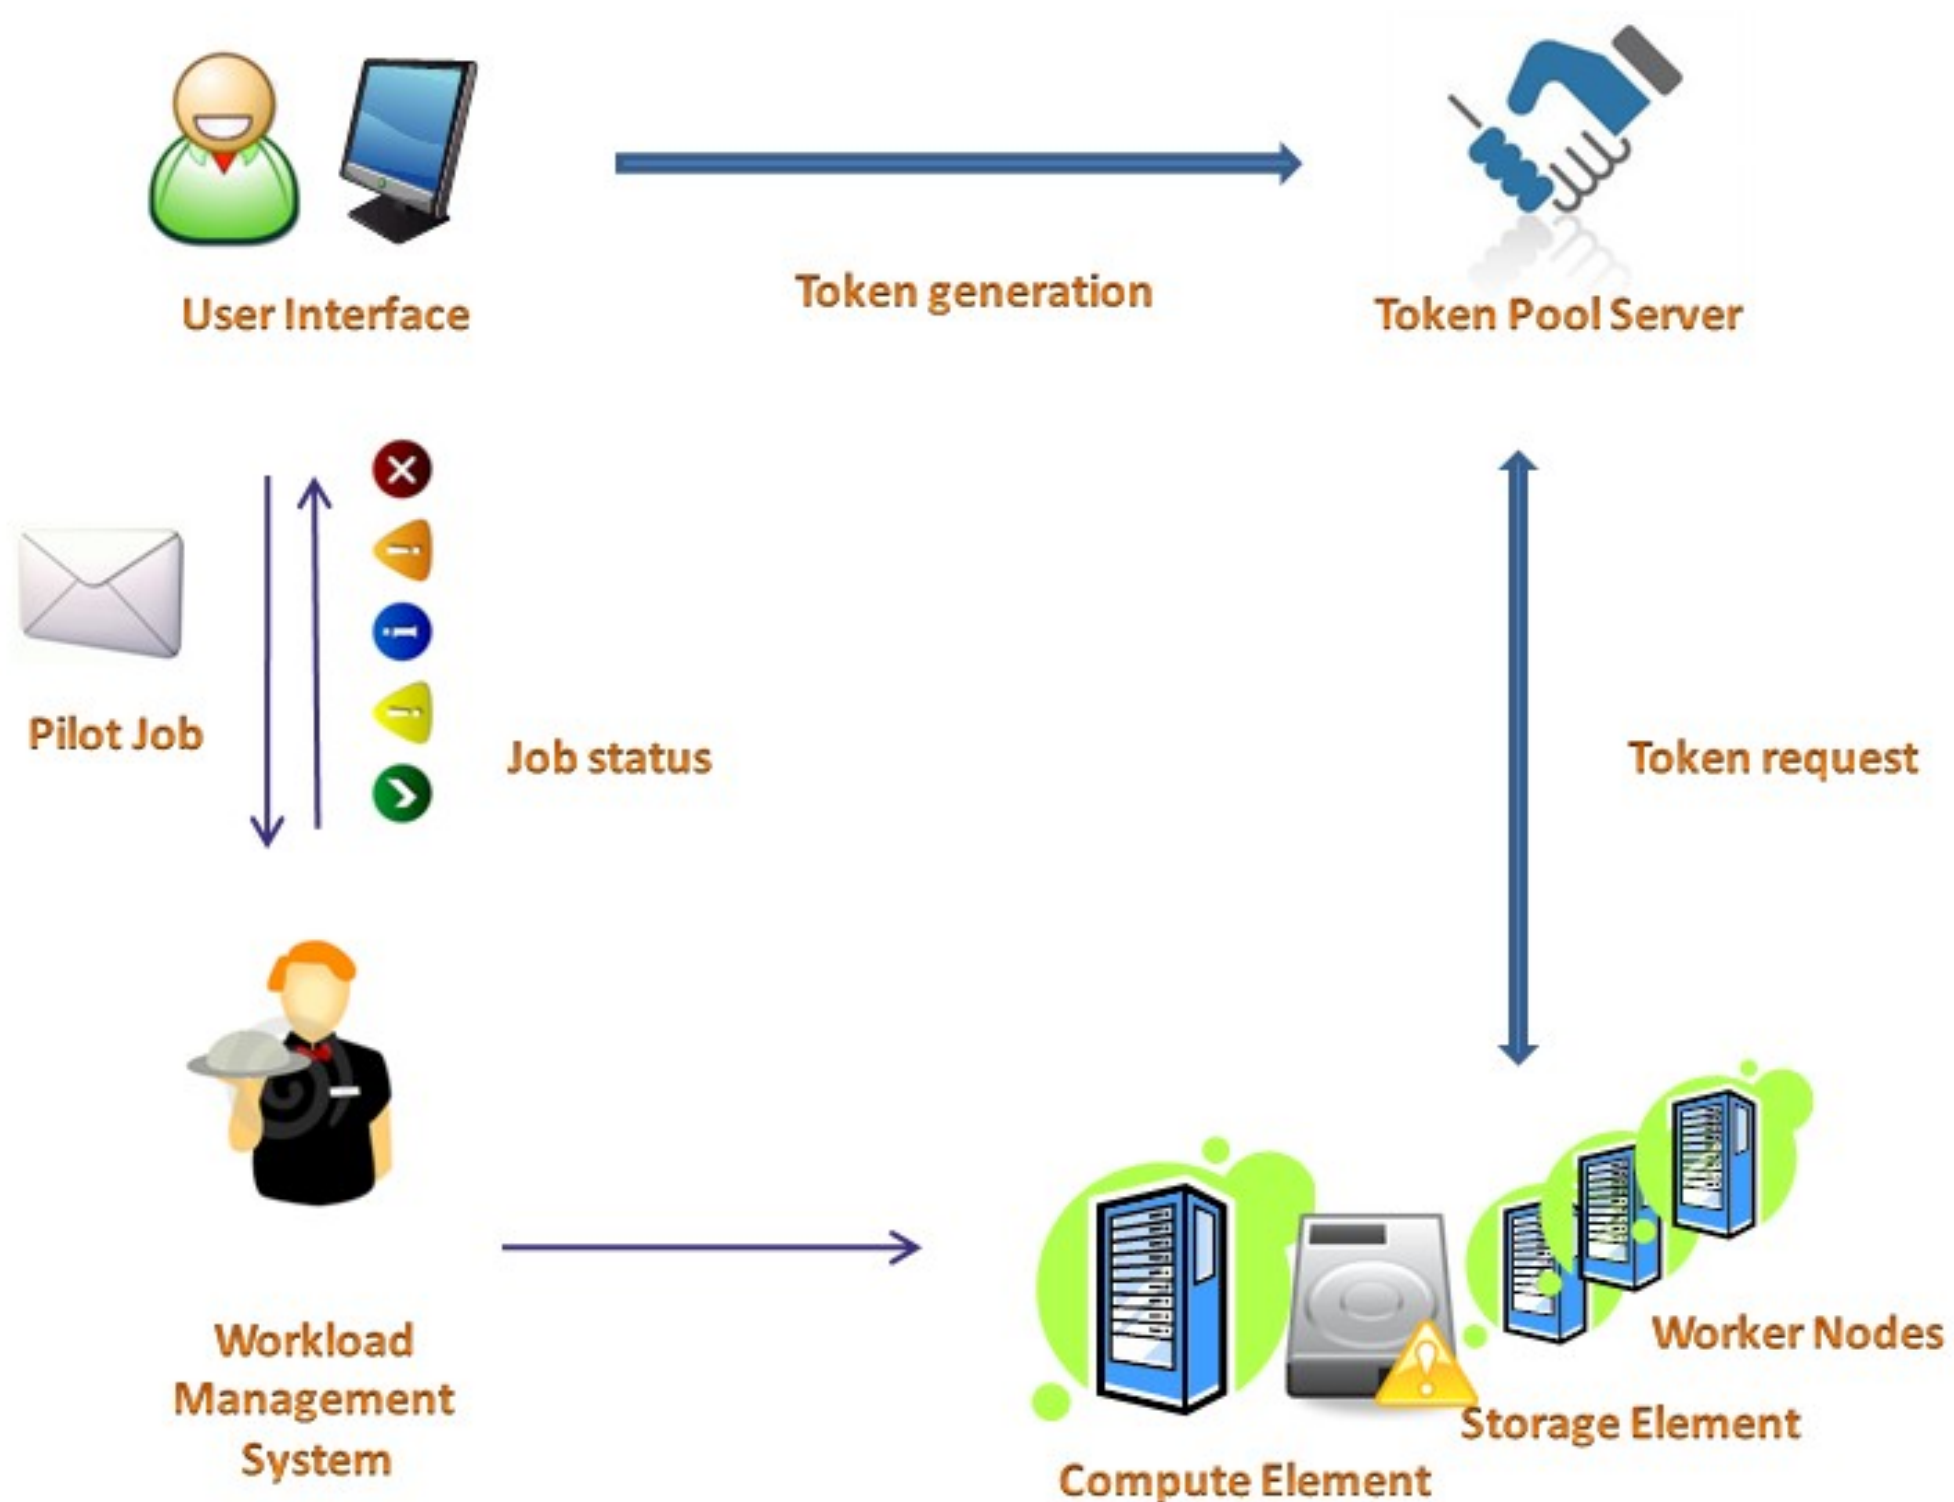
Grid computing

- Grid certificate, you'll be part of a Virtual Organization

Grid computing

- Grid certificate, you'll be part of a Virtual Organization
- (Complex) middleware commands to move data/jobs to computing nodes:
 - Submit job: `glite-wms-job-submit -d $USER -o <jobId> <jdl_file>.jdl`
 - Job status: `glite-wms-job-status -i <jobID>`
 - Copy output to dir: `glite-wms-job-output --dir <dirname> -i <jobID>`
 - Show storage elements: `lcg-infosites --vo ncf se`
 - ls: `lfc-ls -l $LFC_HOME/joaquin`
 - Upload file (copy-register): `lcg-cr --vo ncf -d srm://srm.grid.sara.nl:8443/pnfs/grid.sara.nl/data/ncf/joaquin/<remotename> -l lfn:/grid/ncf/joaquin/<remotename> "file://$PWD/<localname>"`
 - Retrieve file from SE: `lcg-cp --vo ncf lfn:/grid/ncf/joaquin/<remotename> file://$PWD/<localname>`

Token pool servers



Video

