CS 61C: Great Ideas in Computer Architecture

MapReduce
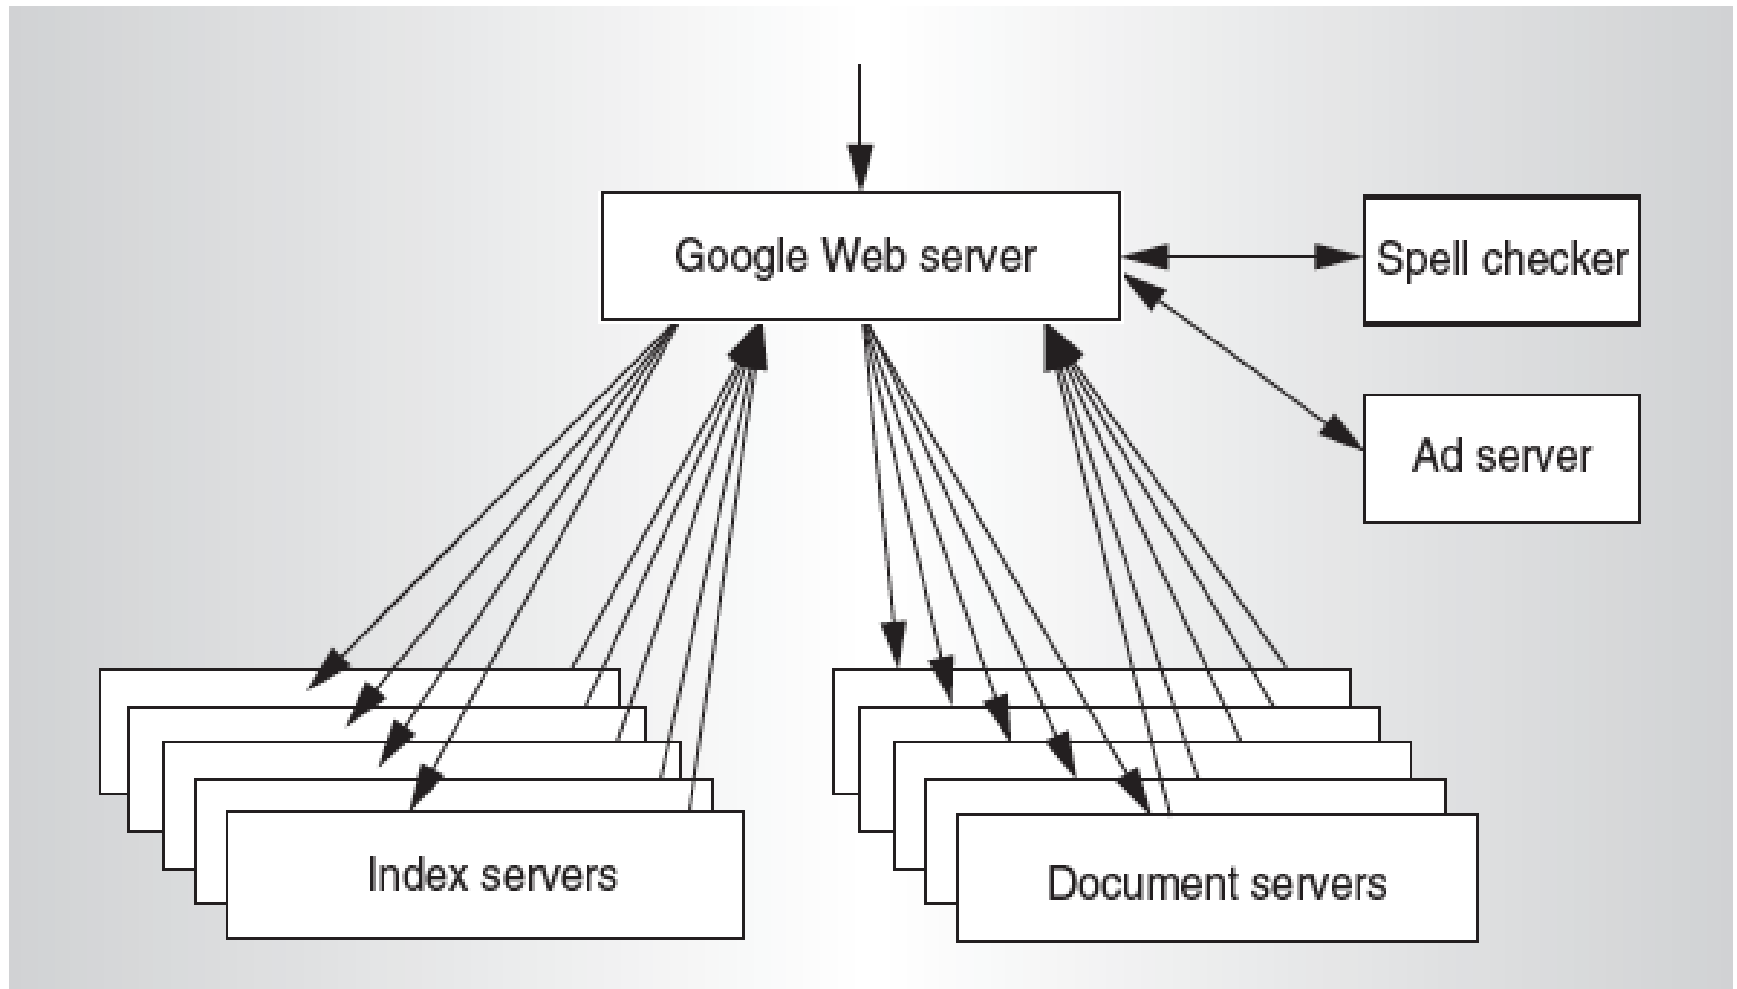
Instructor: Alan Christopher

Review of Last Lecture (1/2)

- Warehouse Scale Computing
  - Example of parallel processing in the post-PC era
  - Servers on a rack, rack part of cluster
  - Issues to handle include load balancing, failures, power usage (sensitive to cost & energy efficiency)
  - PUE = Total building power / IT equipment power

Request-Level Parallelism (RLP) In WSCs

- Hundreds or thousands of requests per sec
  - Many users accessing popular Internet services like web search, social networking, …
  - Such requests are largely independent
    - Often involve read-mostly databases
    - Rarely involve strict read–write data sharing or synchronization across requests
- Computation easily partitioned within a request and across different requests

Google Query-Serving Architecture

Summer 2014 -- Lecture #25

Anatomy of a Web Search

- Google "Leonhard Euler"

Anatomy of a Web Search (1 of 3)

- Google "Leonhard Euler"
    - Direct request to "closest" Google Warehouse Scale Computer
    - Front-end load balancer directs request to one of many arrays (cluster of servers) within WSC
    - Within array, select one of many Google Web Servers (GWS) to handle the request and compose the response pages
    - GWS communicates with Index Servers to find documents that contain the search words, "Leonhard", "Euler", uses location of search as well
    - Return document list with associated relevance score

Anatomy of a Web Search (2 of 3)

- In parallel,
  - Ad system: run ad auction for bidders on search terms
  - Get images of various Leonhard Eulers
- Use docids (document IDs) to access indexed documents
- Compose the page
  - Result document extracts (with keyword in context) ordered by relevance score
  - Sponsored links (along the top) and advertisements (along the sides)

Anatomy of a Web Search (3 of 3)

- Implementation strategy
  - Randomly distribute the entries
  - Make many copies of data (a.k.a. "replicas")
  - Load balance requests across replicas
- Redundant copies of indices and documents
  - Breaks up hot spots, e.g. "Guardians of the Galaxy"
  - Increases opportunities for request-level parallelism
  - Makes the system more tolerant of failures

Agenda

- Data Level Parallelism
- MapReduce
  - Background
  - Design
  - Theory
- Administrivia
- More MapReduce
  - The Combiner + Example 1: Word Count
  - Execution Walkthrough
  - Example 2: PageRank (Time Permitting)

Data-Level Parallelism (DLP)

- Multiple kinds:
  - Lots of data in memory that can be operated on in parallel (e.g. adding together 2 arrays)
  - Lots of data on many disks that can be operated on in parallel (e.g. searching for documents)
1) SIMD does Data-Level Parallelism (DLP) in memory
2) Today's lecture and Lab 12 does DLP across many servers and disks using MapReduce

Agenda

- Data Level Parallelism
- MapReduce
  - Background
  - Design
  - Theory
- Administrivia
- More MapReduce
  - The Combiner + Example 1: Word Count
  - Execution Walkthrough
  - Example 2: PageRank (Time Permitting)

What is MapReduce?

- Simple data-parallel programming model designed for scalability and fault-tolerance

- Pioneered by Google
  - Processes >25 petabytes of data per day

- Popularized by open-source Hadoop project
  - Used at Yahoo!, Facebook, Amazon, …

What is MapReduce used for?

- At Google:
  - Index construction for Google Search
  - Article clustering for Google News
  - Statistical machine translation
  - For computing multi-layer street maps
- At Yahoo!:
  - "Web map" powering Yahoo! Search
  - Spam detection for Yahoo! Mail
- At Facebook:
  - Data mining
  - Ad optimization
  - Spam detection

MapReduce Design Goals

1. Scalability to large data volumes:
    - 1000's of machines, 10,000's of disks

2. Cost-efficiency:
    - Commodity machines (cheap, but unreliable)
    - Commodity network
    - Automatic fault-tolerance (fewer administrators)
    - Easy to use (fewer programmers)

Jeffrey Dean and Sanjay Ghemawat, "MapReduce: Simplified Data Processing on Large Clusters," *6th USENIX Symposium on Operating Systems Design and Implementation, 2004*. (optional reading, linked on course homepage – a digestible CS paper at the 61C level)

MapReduce Processing: "Divide and Conquer" (1/3)

- Apply Map function to user supplied record of key/value pairs
  - Slice data into "shards" or "splits" and distribute to workers
  - Compute set of intermediate key/value pairs

```
map(in_key,in_val):
      // DO WORK HERE
      emit(interm_key,interm_val)
```

MapReduce Processing: "Divide and Conquer" (2/3)

- Apply Reduce operation to all values that share same key in order to combine derived data properly
  - Combines all intermediate values for a particular key
  - Produces a set of merged output values

```
reduce(interm_key,list(interm_val)):
        // DO WORK HERE
        emit(out_key, out_val)
```

MapReduce Processing: "Divide and Conquer" (3/3)

- User supplies Map and Reduce operations in functional model
  - Focus on problem, let MapReduce library deal with messy details
  - Parallelization handled by framework/library
  - Fault tolerance via re-execution

Typical Hadoop Cluster

Aggregation switch

Rack switch

8 gigabit
1 gigabit

Node
Disks

Node
Disks

Node
Disks

Node
Disks

Node
Disks
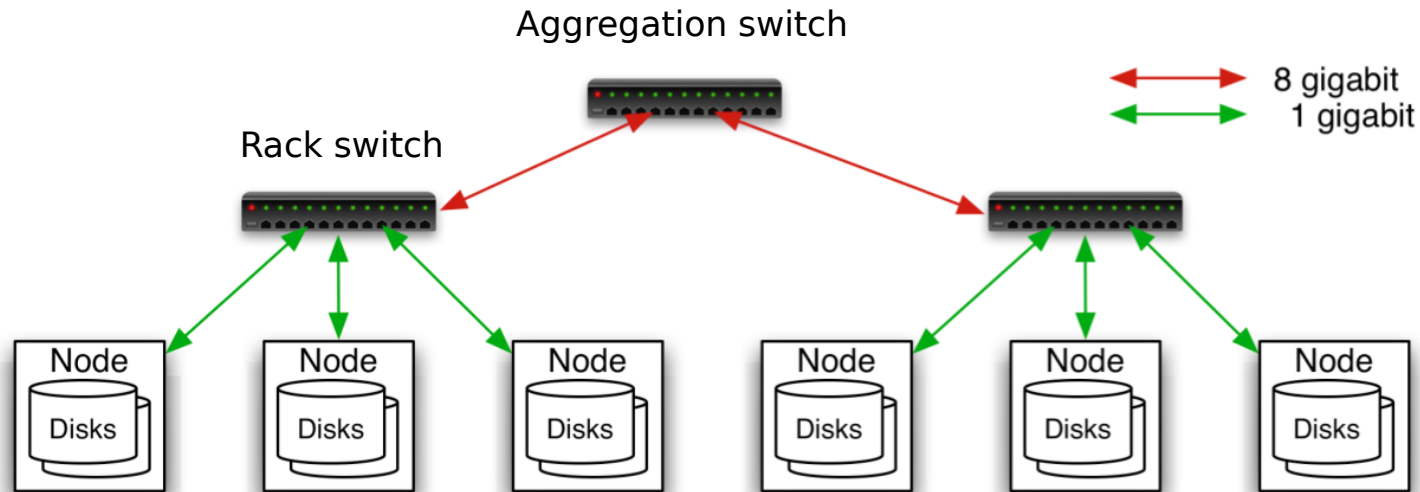
Node
Disks

- 40 nodes/rack, 1000-4000 nodes in cluster
- 1 Gbps bandwidth within rack, 8 Gbps out of rack
- Node specs (Yahoo terasort):
  8 x 2GHz cores, 8 GB RAM, 4 disks (= 4 TB?)

Agenda

- Data Level Parallelism
- MapReduce
  - Background
  - Design
  - Theory
- Administrivia
- More MapReduce
  - The Combiner + Execution Walkthrough
  - Example 1: Word Count
  - Example 2: PageRank (Time Permitting)

Administrivia

- Project 3 (partners) due Sunday 8/10
- Final Review – Sat 08/09, 2-5pm in 2060 VLSB
- Final – Fri 8/15, 9am-12pm, 155 Dwinelle
  - MIPS Green Sheet provided again
  - Two-sided handwritten cheat sheet
    - Can use the back side of your midterm cheat sheet!

Agenda

- Data Level Parallelism
- MapReduce
  - Background
  - Design
  - Theory
- Administrivia
- More MapReduce
  - The Combiner + Example 1: Word Count
  - Execution Walkthrough
  - Example 2: PageRank (Time Permitting)

The Combiner (Optional)

- One missing piece for our first example:
    - Many times, the output of a single mapper can be "compressed" to save on bandwidth and to distribute work (usually more map tasks than reduce tasks)
    - To implement this, we have the combiner:

```
combiner(interm_key,list(interm_val)):
        // DO WORK (usually like reducer)
        emit(interm_key2, interm_val2)
```

Our Final Execution Sequence

- <u>Map</u> – Apply operations to all input key, val

- <u>Combine</u> – Apply reducer operation, but distributed across map tasks

- <u>Reduce</u> – Combine all values of a key to produce desired output

MapReduce Processing Example: Count Word Occurrences (1/2)

- Pseudo Code: for each word in input, generate <key=word, value=1>
- Reduce sums all counts emitted for a particular word across all mappers

```
map(String input_key, String input_value):
  // input_key: document name
  // input_value: document contents
  for each word w in input_value:
    EmitIntermediate(w, "1"); // Produce count of words
 combiner: (same as below reducer)
 reduce(String output_key, Iterator intermediate_values):
  // output_key: a word
  // intermediate_values: a list of counts
  int result = 0;
  for each v in intermediate_values:
    result += ParseInt(v); // get integer from key-value
  Emit(output_key, result);
```

MapReduce Processing Example: Count Word Occurrences (2/2)

Distribute

that that is	is that that	is not is not	is that it it is
Map 1	Map 2	Map 3	Map 4
is 1, that 2	is 1, that 2	is 2, not 2	is 2, it 2, that 1

Shuffle

is 1,1,2,2
it 2

Reduce 1

is 6; it 2

that 2,2,1
not 2
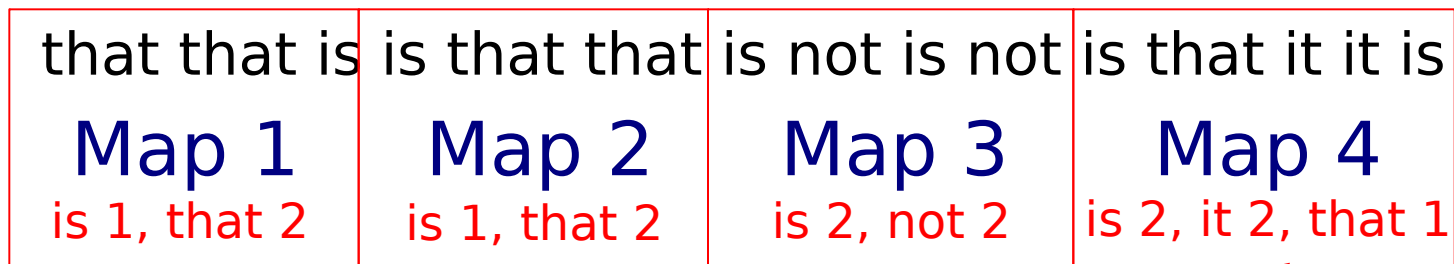
Reduce 2

not 2; that 5

Collect

is 6; it 2; not 2; that 5

Agenda

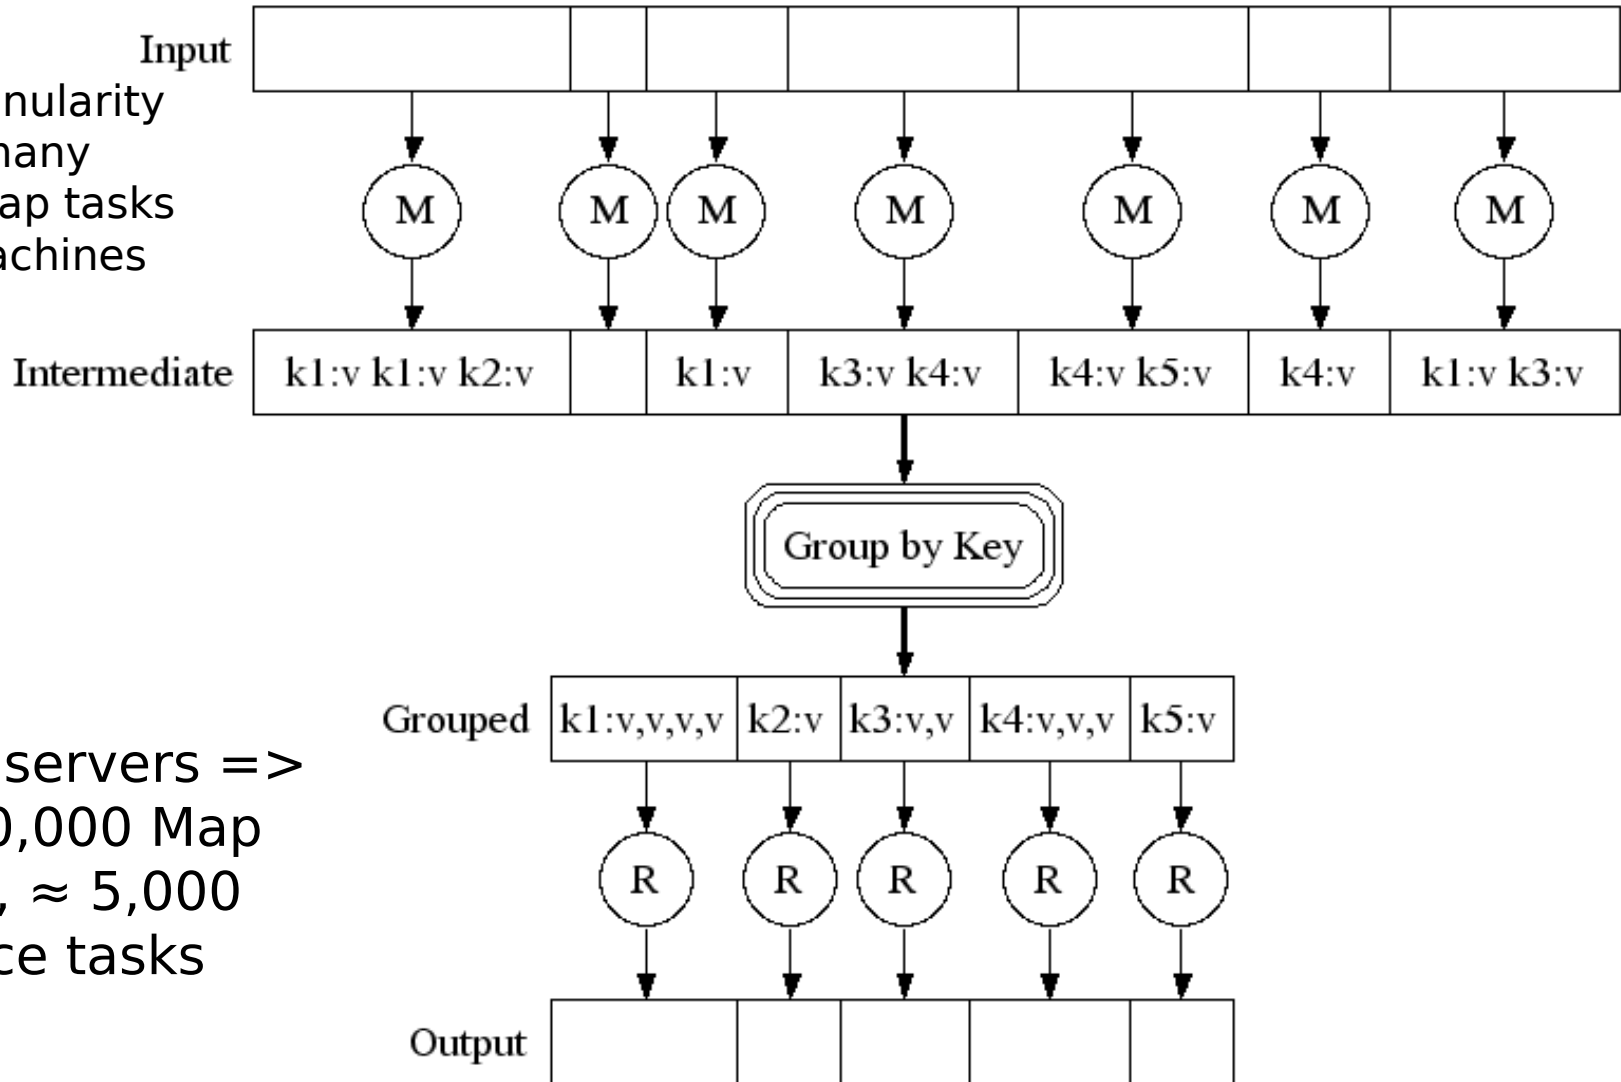- Data Level Parallelism
- MapReduce
  - Background
  - Design
  - Theory
- Administrivia
- More MapReduce
  - The Combiner + Example 1: Word Count
  - Execution Walkthrough
  - Example 2: PageRank (Time Permitting)

Execution Setup

- Map invocations distributed by partitioning input data into M *splits*
  - Typically 16 MB to 64 MB per piece
- Input processed in parallel on different servers
- Reduce invocations distributed by partitioning intermediate key space into R pieces
  - e.g. hash(key) mod R
- User picks M >> # servers, R > # servers
  - Big M helps with load balancing, recovery from failure
  - One output file per R invocation, so not too many

MapReduce Execution

Fine granularity tasks: many more map tasks than machines

Input							

M   M M   M   M   M   M

Intermediate	k1:v k1:v k2:v		k1:v	k3:v k4:v	k4:v k5:v	k4:v	k1:v k3:v

Group by Key

Grouped	k1:v,v,v,v	k2:v	k3:v,v	k4:v,v,v	k5:v

2000 servers =>
≈ 200,000 Map
Tasks, ≈ 5,000
Reduce tasks

R   R   R   R   R

Output					

MapReduce Processing

Su Shuffle phase #26
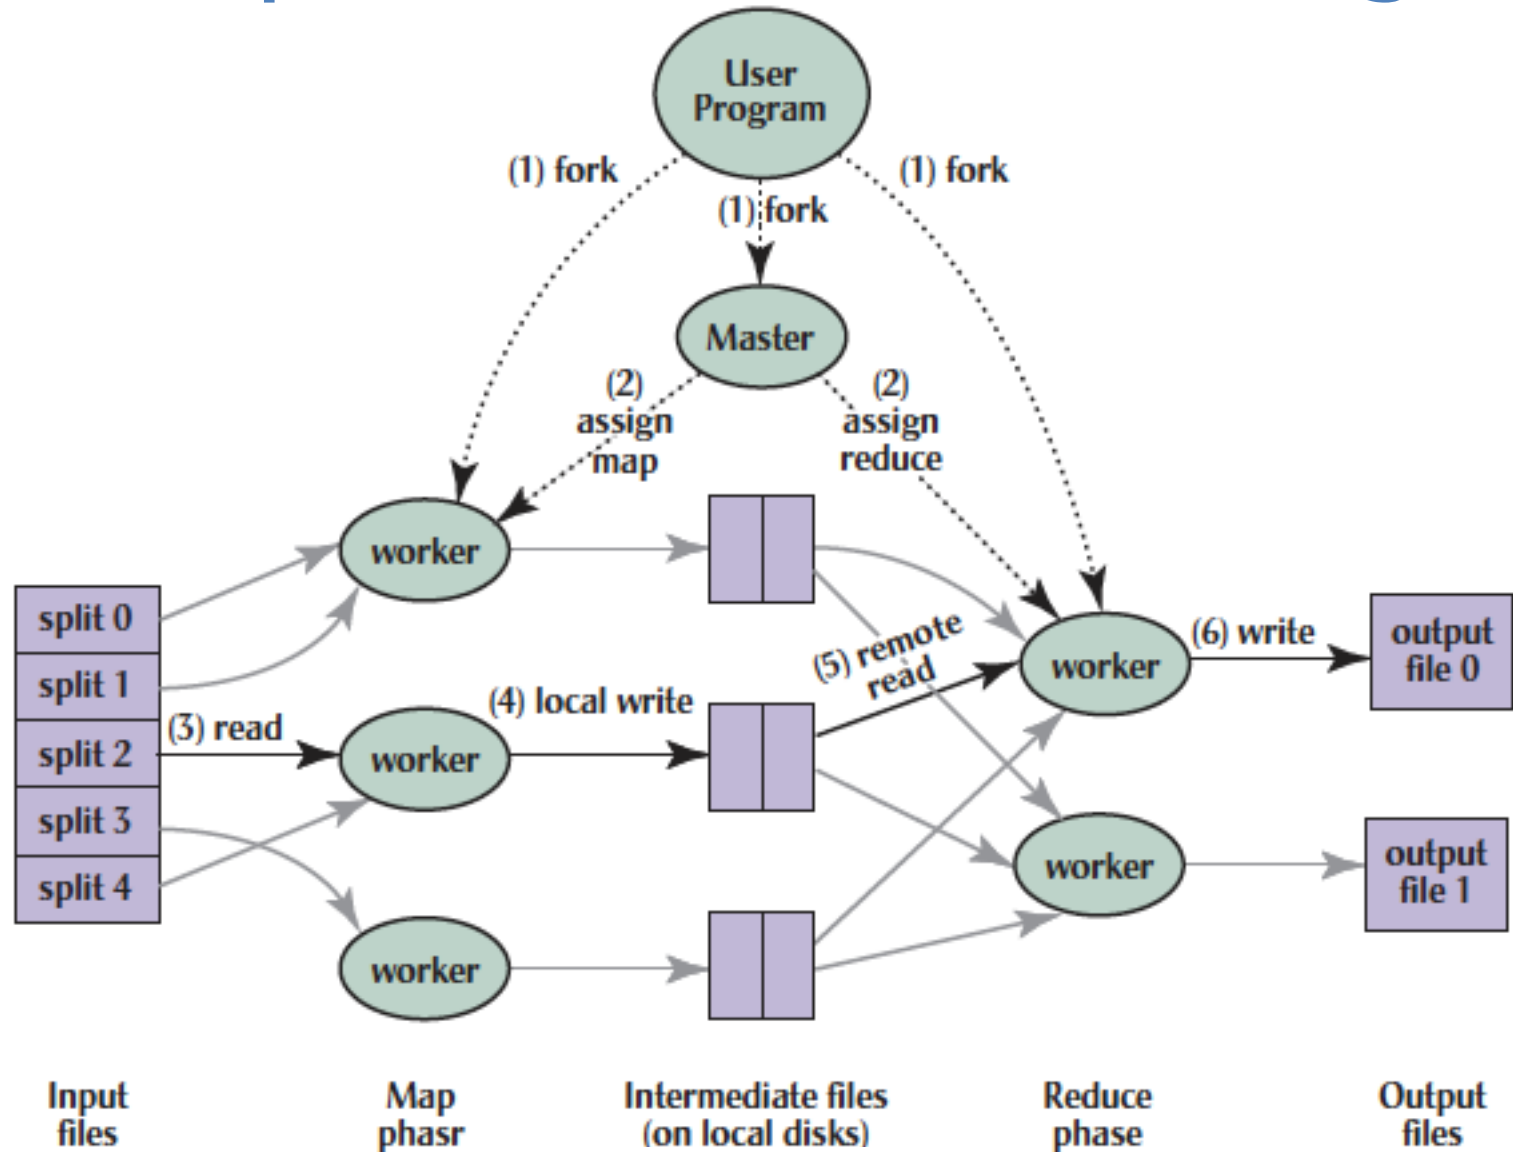
MapReduce Processing

1. MR 1st splits the input files into *M* "splits" then starts many copies of program on servers

1. MR 1st splits the input files into *M* "splits" then starts many copies of program on servers

Input files | Map phasr | Intermediate files (on local disks) | Reduce phase | Output files

Shuffle phase
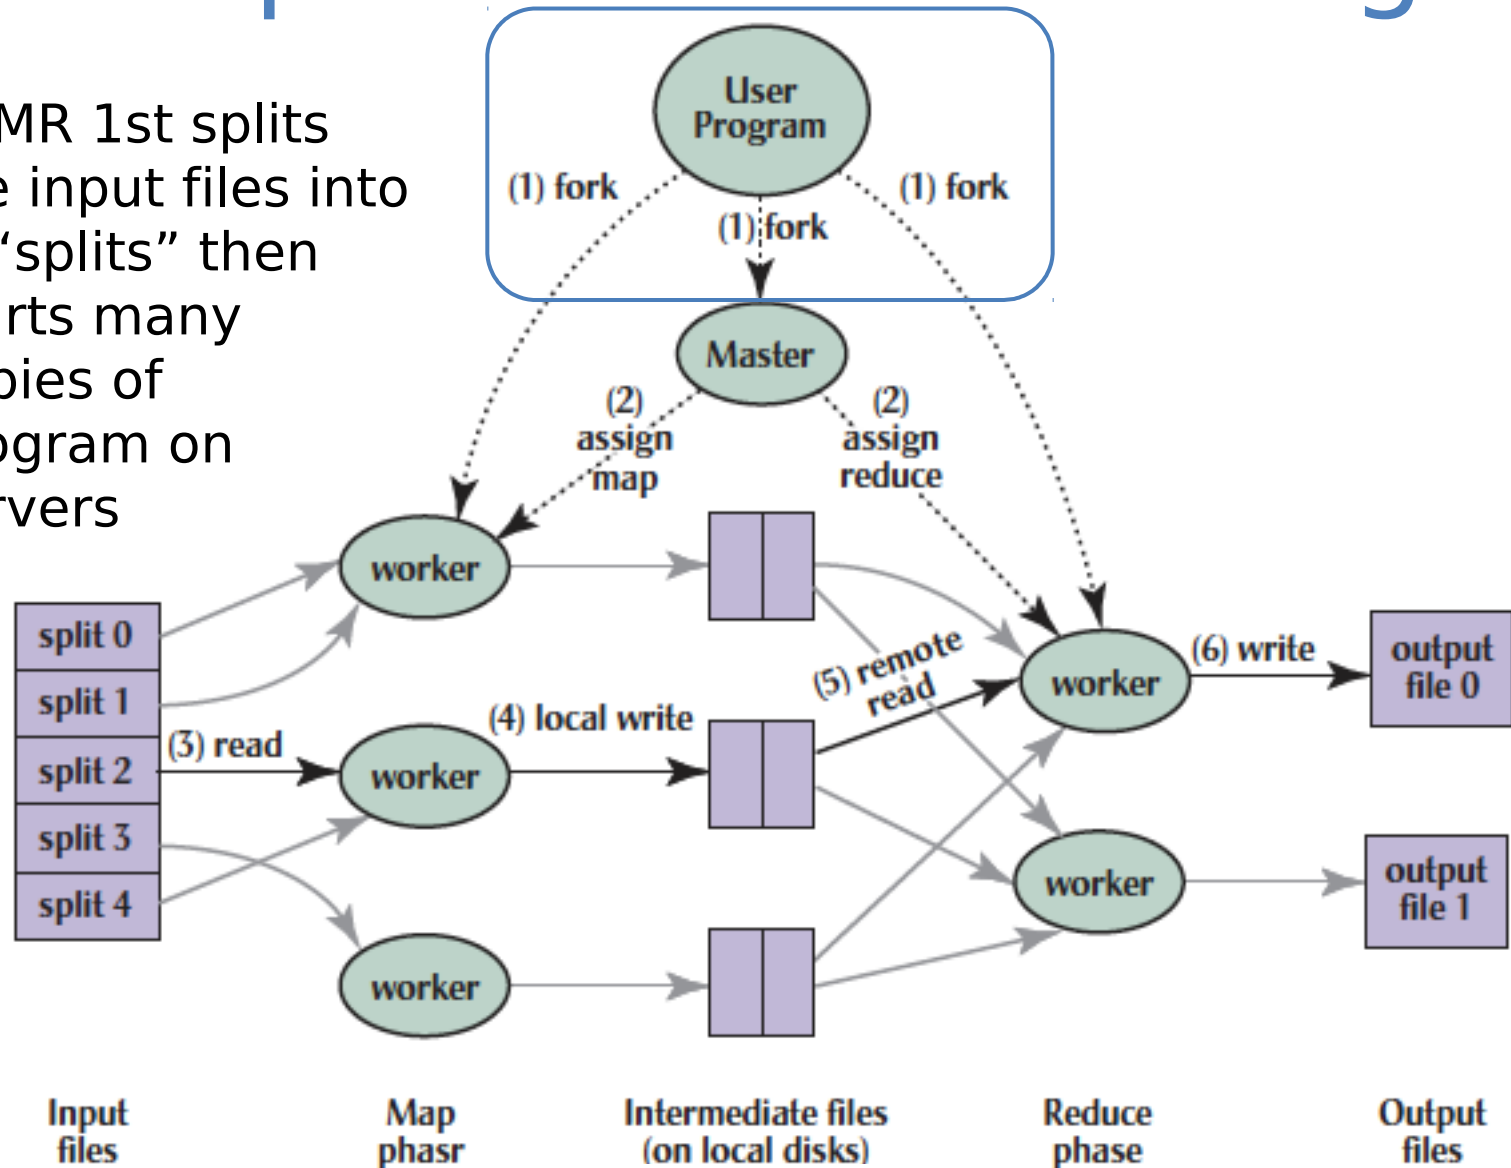
MapReduce Processing

2. One copy (the master) is special. The rest are workers. The master picks idle workers and assigns each 1 of M map tasks or 1 of R reduce tasks.

Su Shuffle phase #26
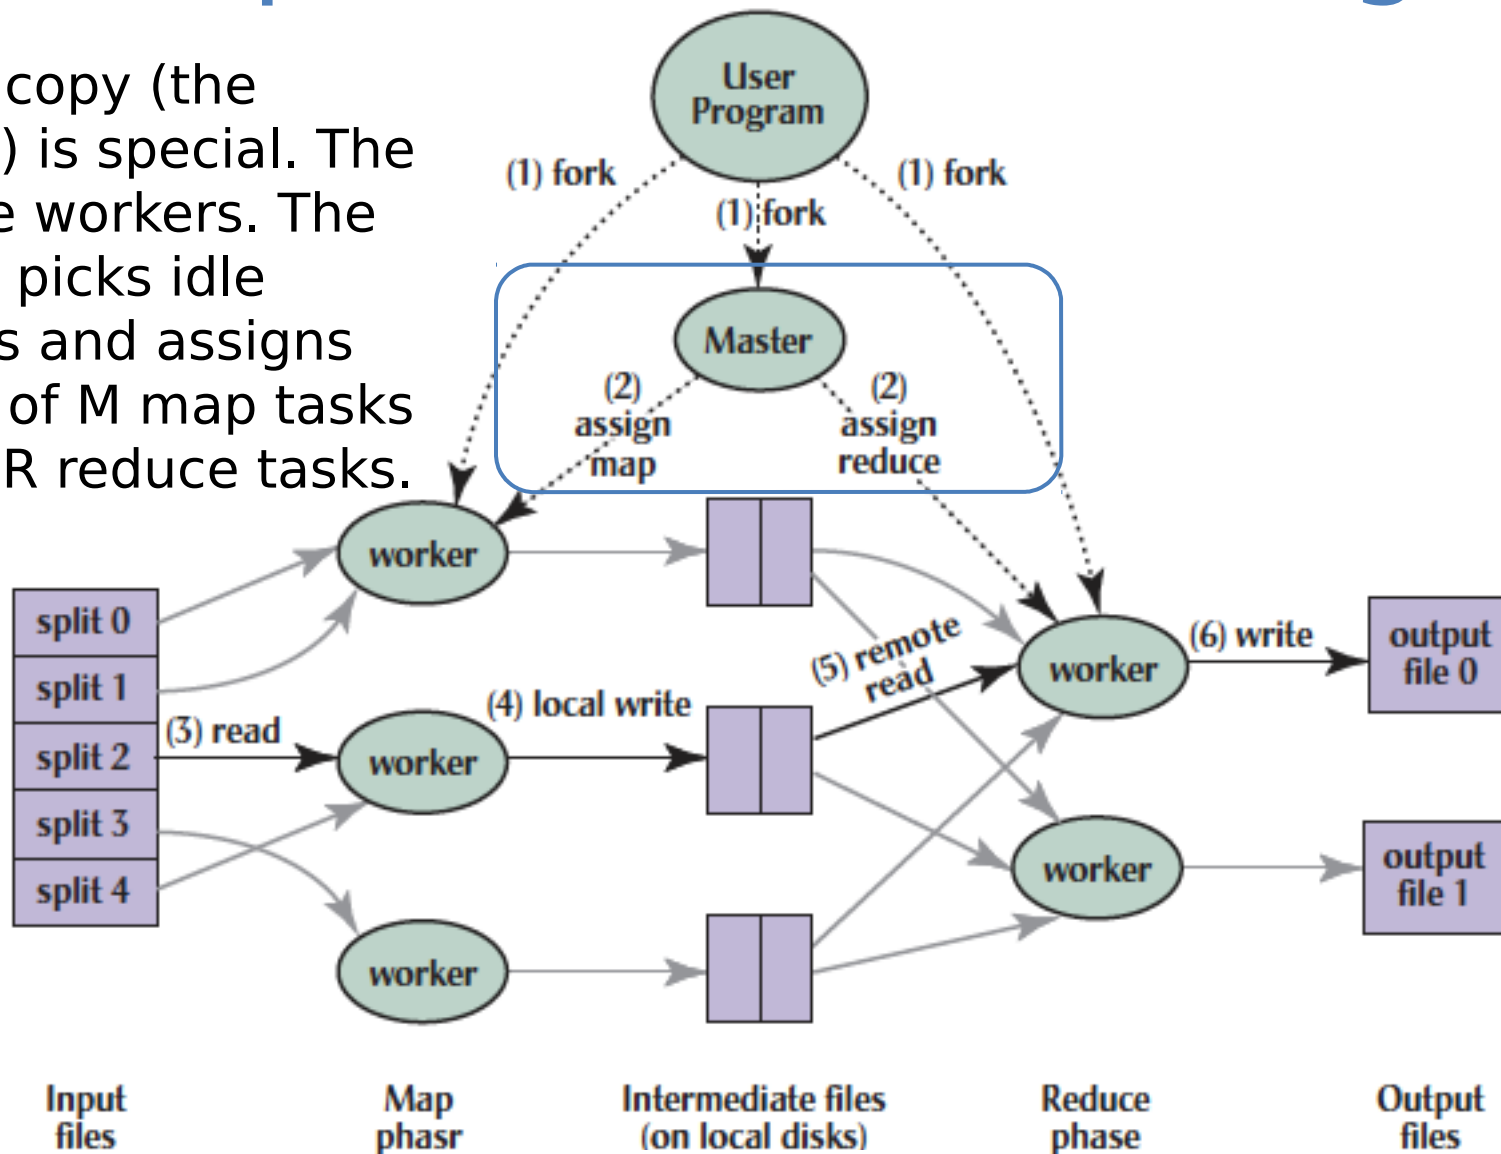
MapReduce Processing

3. A map worker reads the input split. It parses key/value pairs of the input data and passes each pair to the user-defined map function.

(The intermediate key/value pairs produced by the map function are buffered in memory.)

Shuffle phase
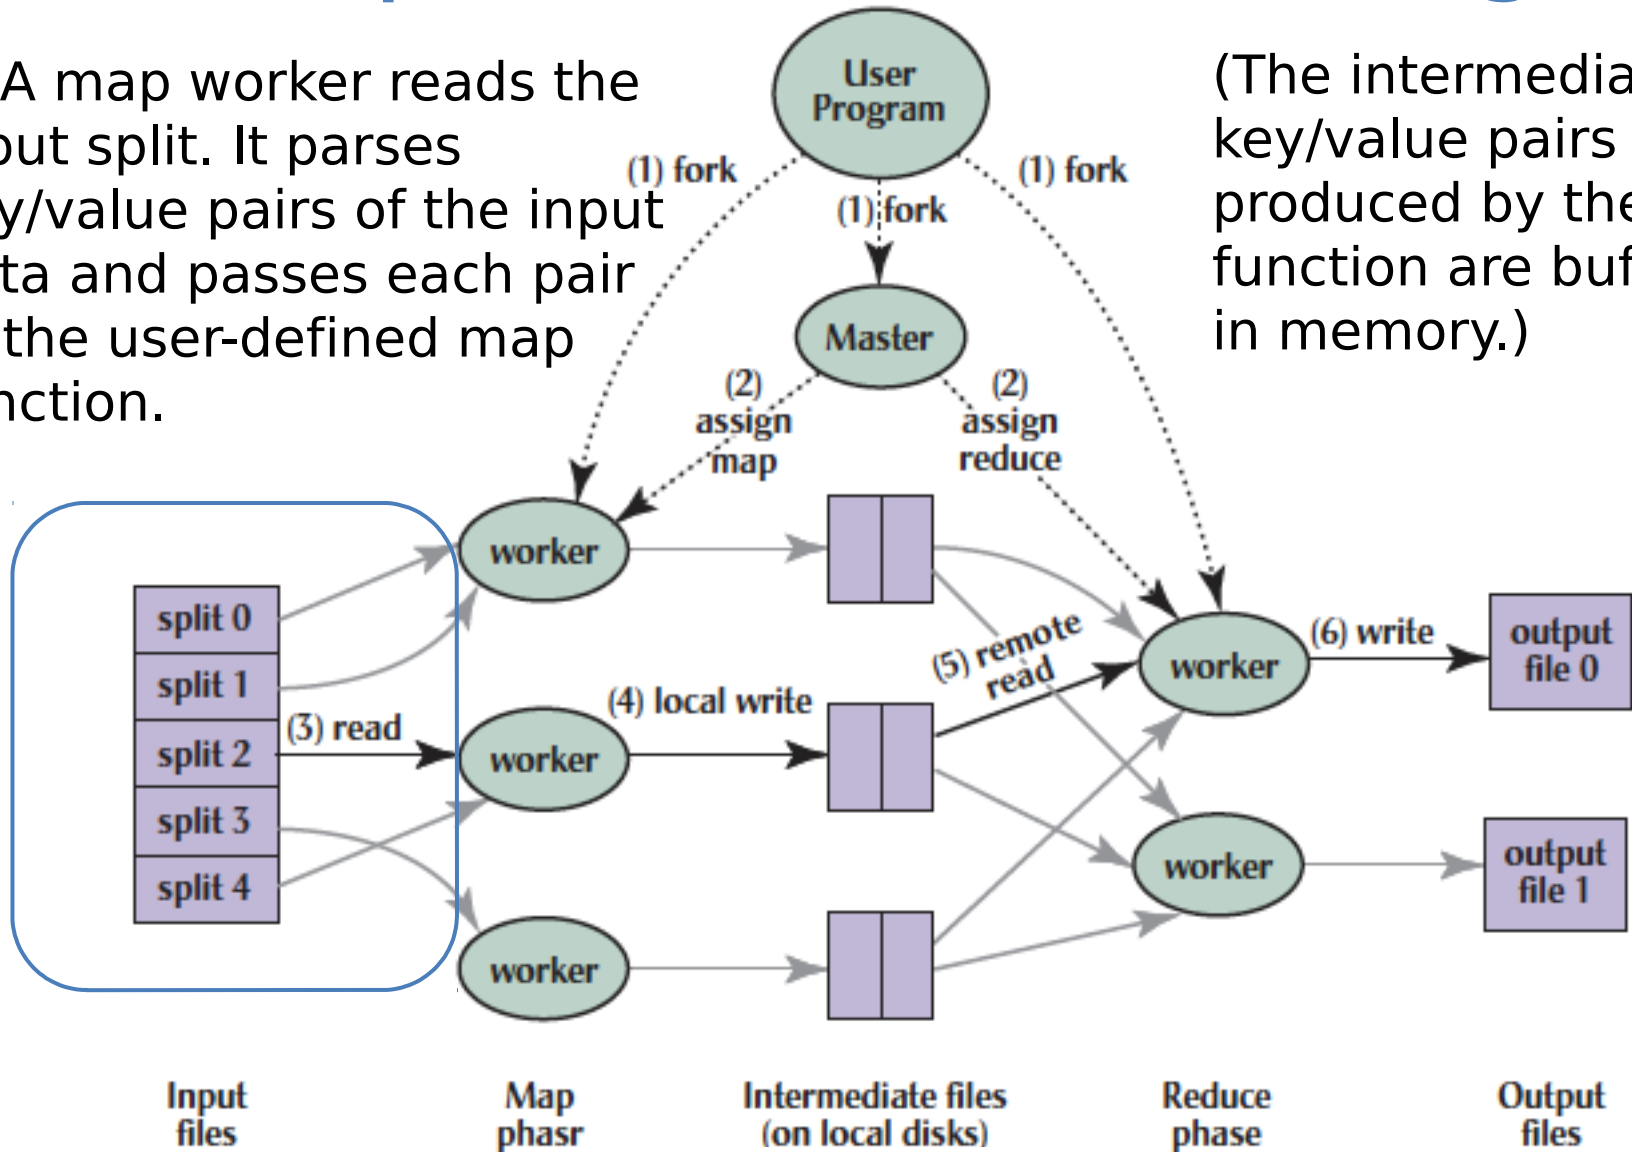
MapReduce Processing

4. Periodically, the buffered pairs are written to local disk, partitioned into R regions by the partitioning function.

Shuffle phase
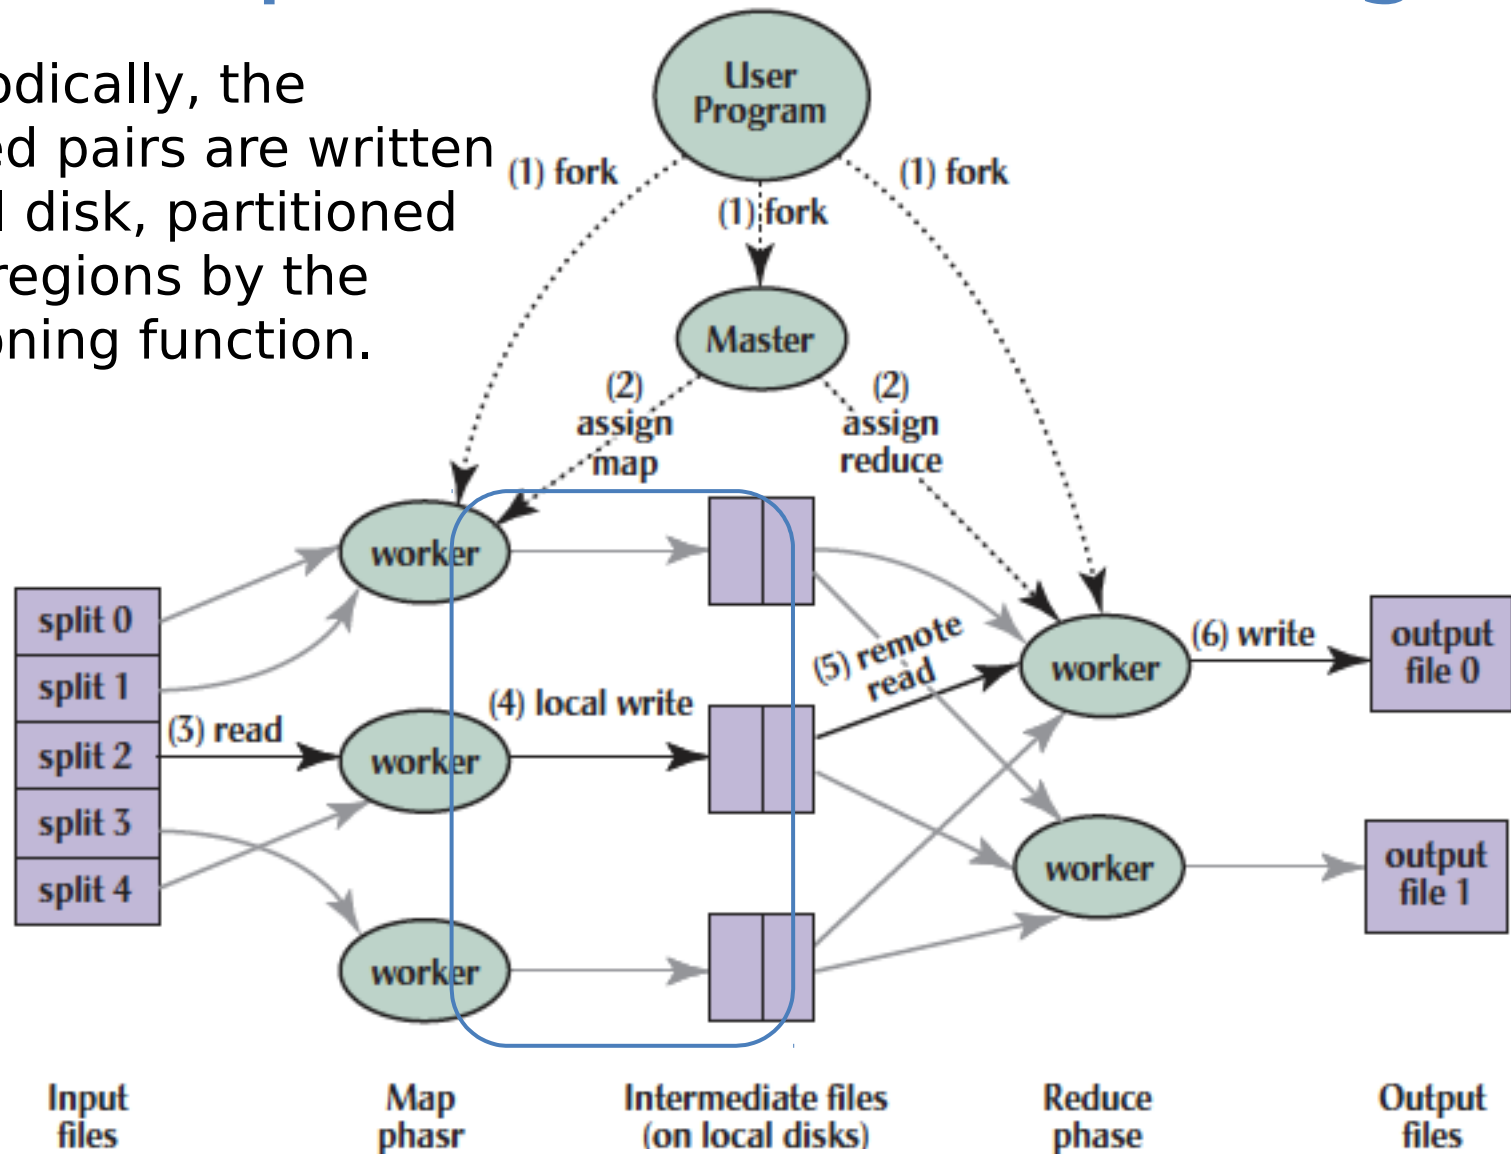
MapReduce Processing

5. When a reduce worker has read all intermediate data for its partition, it sorts it by the intermediate keys so that all occurrences of the same key are grouped together.

(The sorting is needed because typically many different keys map to the same reduce task)

Input files · Map phasr · Intermediate files (on local disks) · Reduce phase · Output files

Shuffle phase
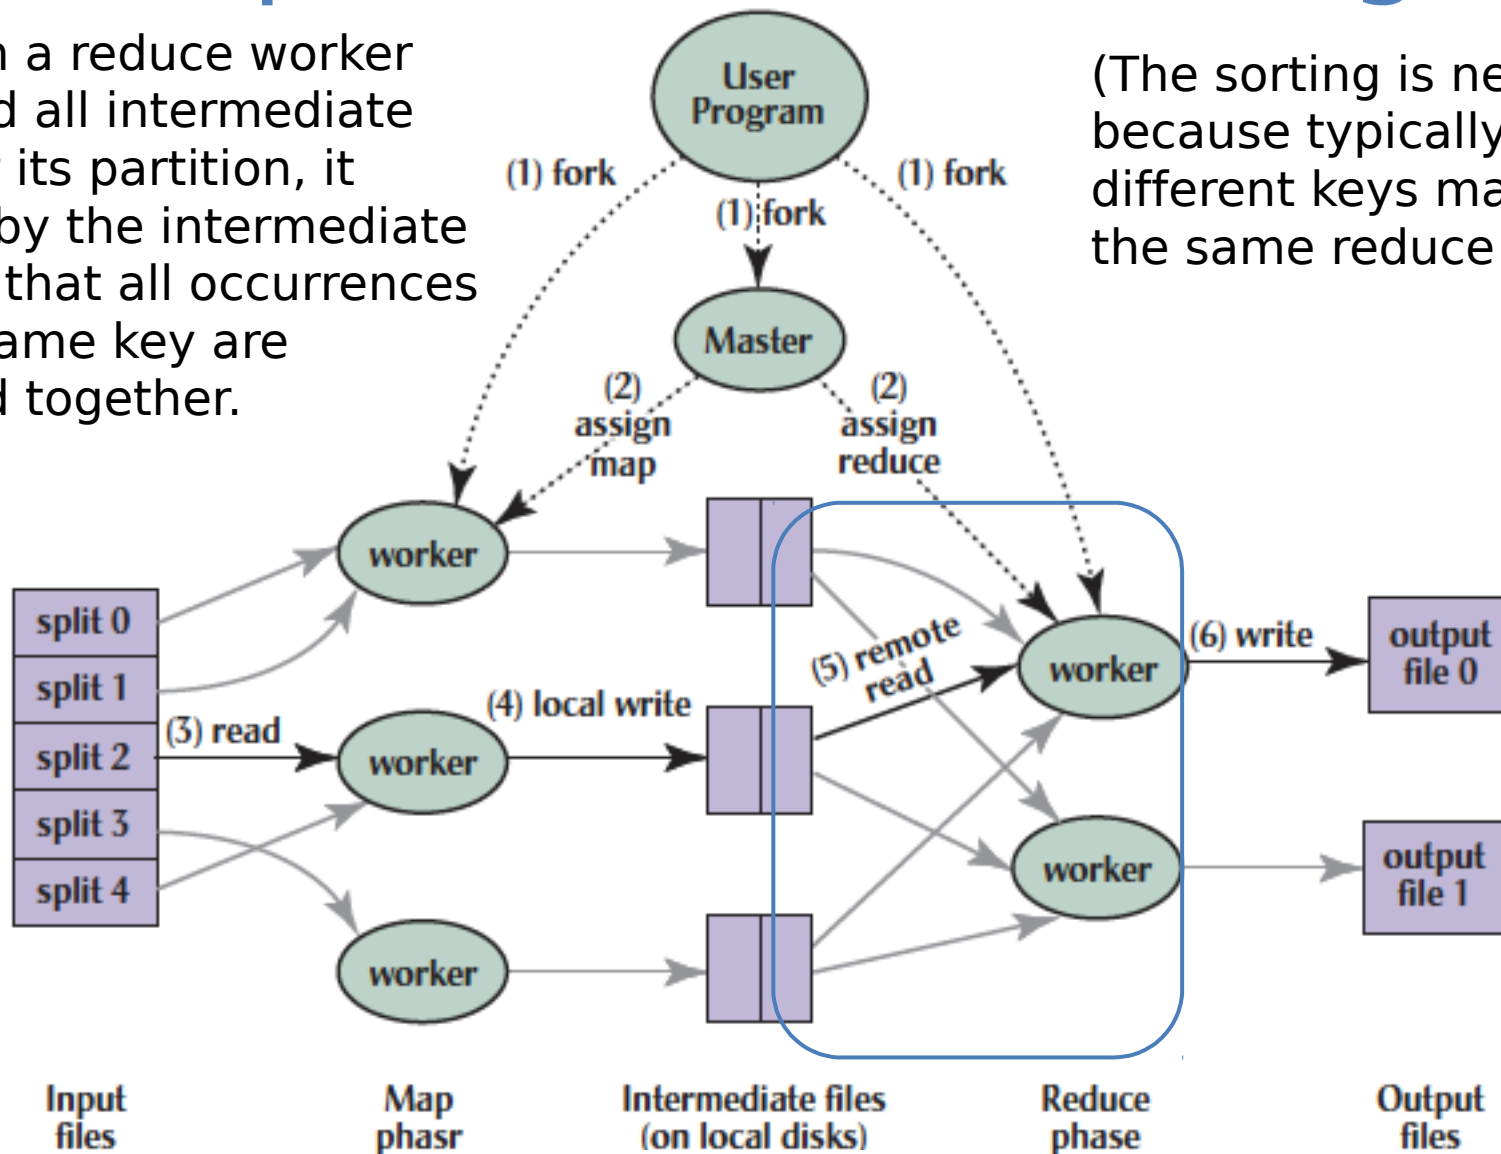
MapReduce Processing

6. Reduce worker iterates over sorted intermediate data and for each unique intermediate key, it passes key and corresponding set of values to the user's reduce function.

The output of the reduce function is appended to a final output file for this reduce partition.

Shuffle phase
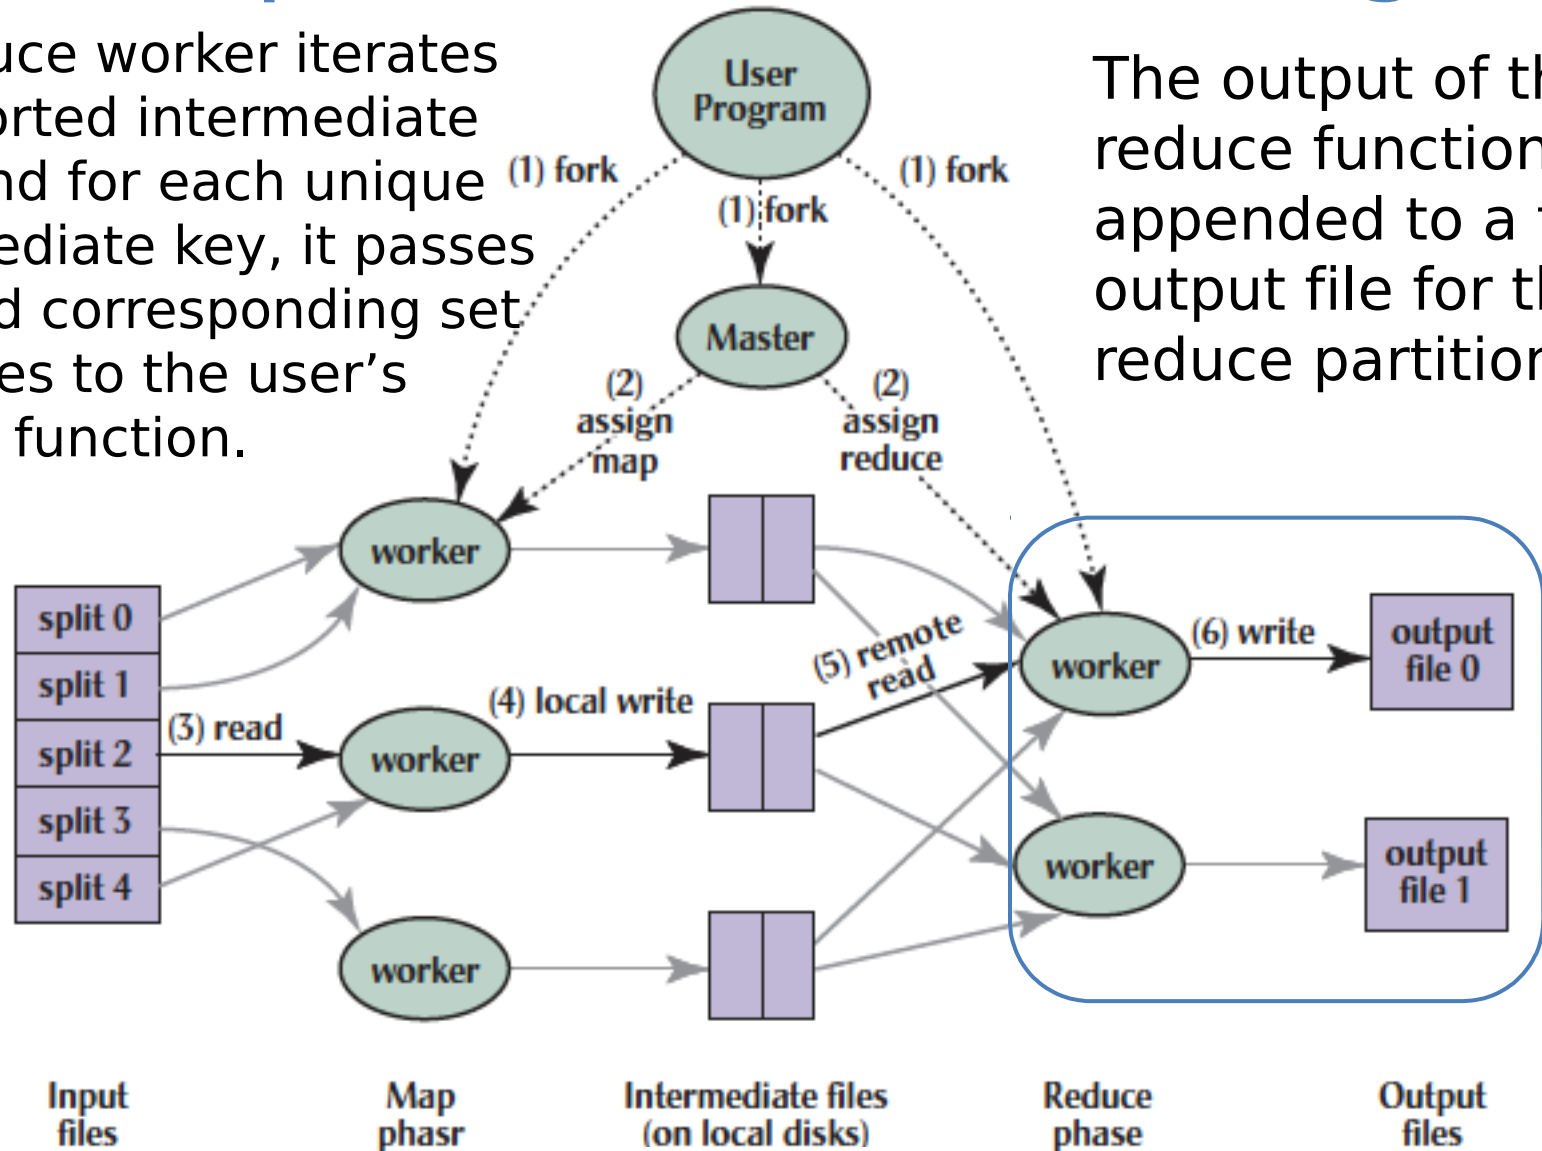
MapReduce Processing

7. When all map and reduce tasks have been completed, the master wakes up the user program. The MapReduce call in user program returns back to user code.

Output of MR is in *R* output files (1 per reduce task, with file names specified by user); often passed into another MR job.

User Program

(1) fork   (1) fork
(1) fork

Master

(2) assign map   (2) assign reduce

worker

split 0
split 1
split 2
split 3
split 4

(3) read

worker

(4) local write

(5) remote read

worker

(6) write

output file 0

worker

output file 1

worker

Input files

Map phasr

Intermediate files (on local disks)

Reduce phase

Output files

Shuffle phase
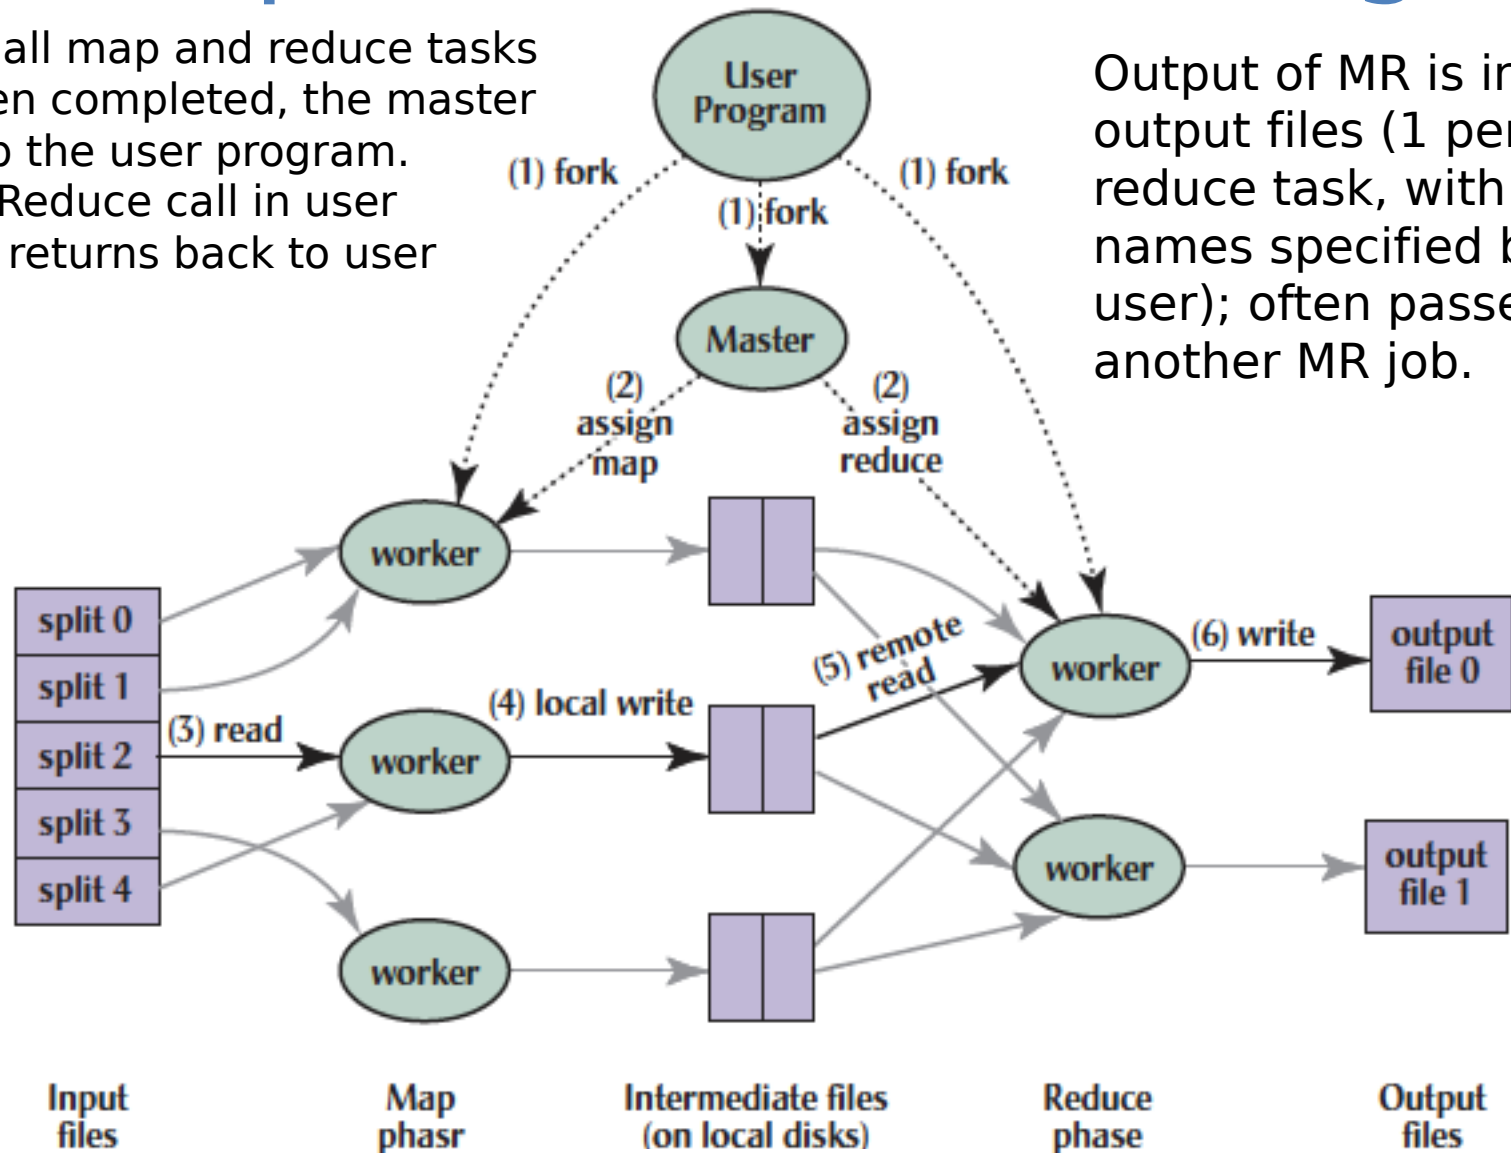
Technology Break

What Does the Master Do?

- For each map task and reduce task, keep track:
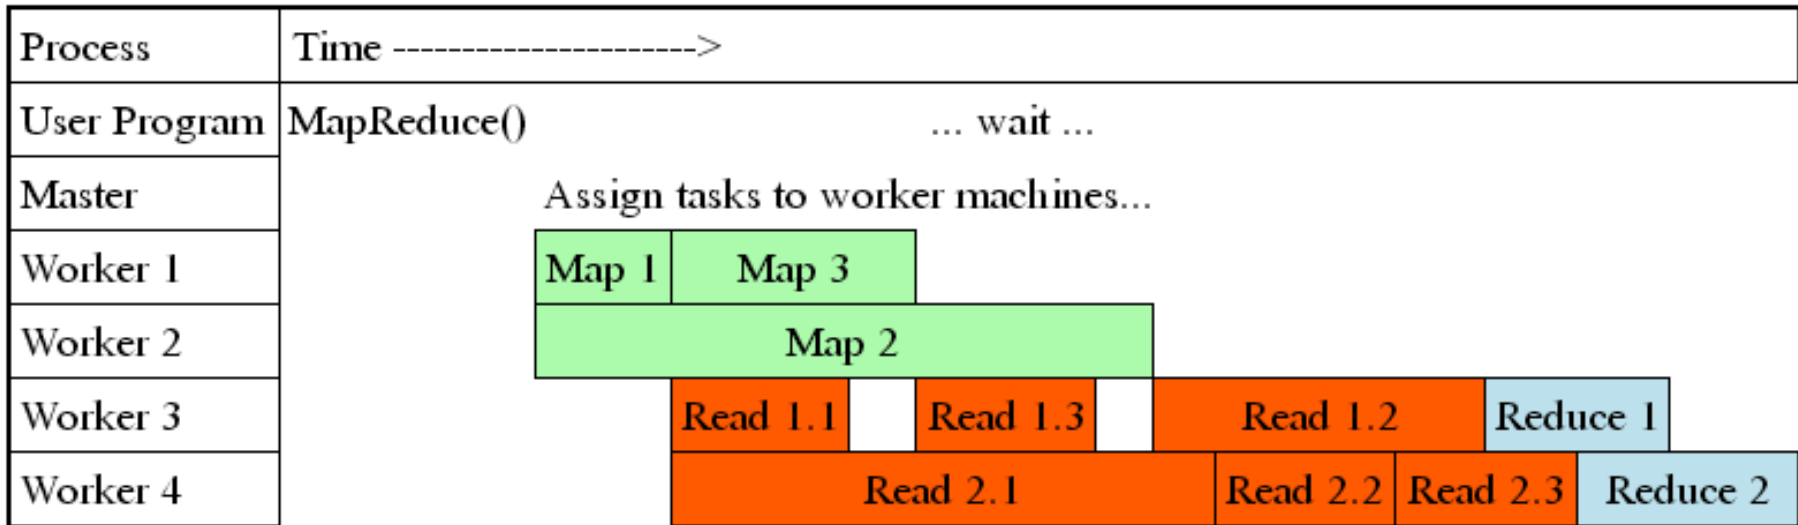 - State: idle, in-progress, or completed
  - Identity of worker server (if not idle)
- For each completed map task
  - Stores location and size of R intermediate files
  - Updates files and size as corresponding map tasks complete
- Location and size are pushed incrementally to workers that have in-progress reduce tasks

MapReduce Processing Time Line

Process	Time --------------------->			
User Program	MapReduce()		... wait ...	
Master		Assign tasks to worker machines...		
Worker 1		Map 1	Map 3	
Worker 2		Map 2		
Worker 3		Read 1.1	Read 1.3	Read 1.2	Reduce 1
Worker 4		Read 2.1	Read 2.2	Read 2.3	Reduce 2

- Master assigns map + reduce tasks to "worker" servers
- As soon as a map task finishes, worker server can be assigned a new map or reduce task
- Data shuffle begins as soon as a given Map finishes
- Reduce task begins as soon as all data shuffles finish
- To tolerate faults, reassign task if a worker server "dies"

MapReduce Failure Handling

- On worker failure:
  - Detect failure via periodic heartbeats
  - Re-execute completed and in-progress map tasks
  - Re-execute in progress reduce tasks
  - Task completion committed through master
- Master failure:
  - Protocols exist to handle (master failure unlikely), but not covered here
- Robust: lost 1600 of 1800 machines once, but finished fine

MapReduce Redundant Execution

- Slow workers significantly lengthen completion time
  - Other jobs consuming resources on machine
  - Bad disks with soft errors transfer data very slowly
  - Weird things: processor caches disabled (!!)
- Solution: Near end of phase, spawn backup copies of tasks
  - Whichever one finishes first "wins"
- Effect: Dramatically shortens job completion time
  - 3% more resources, large tasks 30% faster

Question: Which statement is **NOT** true about about MapReduce?

(B) MapReduce divides computers into 1 master and N-1 workers; masters assigns MR tasks

(G) Towards the end, the master assigns uncompleted tasks again; 1st to finish wins

(P) Reducers can start reducing as soon as they start to receive Map data

(Y) **Reduce worker sorts by intermediate keys to group all occurrences of same key**

Agenda

- Data Level Parallelism
- MapReduce
  - Background
  - Design
  - Theory
- Administrivia
- More MapReduce
  - The Combiner + Example 1: Word Count
  - Execution Walkthrough
  - Example 2: PageRank (Time Permitting)

PageRank: How Google Search Works

- Earlier: RLP – how Google handles searching its huge index
- Now: How does Google generate that index?
- PageRank is the famous algorithm behind the "quality" of Google's results
  - Uses link structure to rank pages, instead of matching only against content (keyword)
  - Modern google search more sophisticated, but PageRank still plays a role
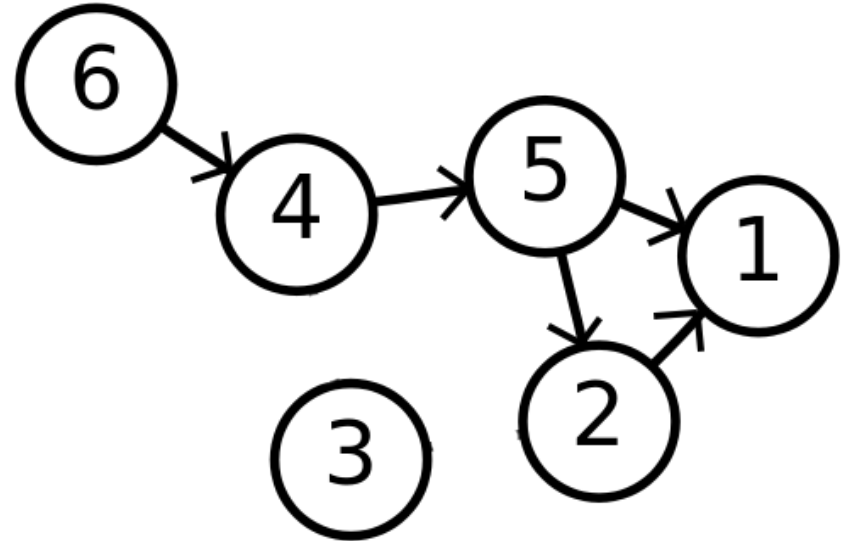
A Quick Review of CS Theory: Graphs

- <u>Def</u>: A set of objects connected by links
- The "objects" are called Nodes
- The "links" are called Edges
- Nodes: {1, 2, 3, 4, 5, 6}
- Edges:
  {(6, 4), (4, 5), (4, 3), (3, 2), (5, 2), (5, 1), (1, 2)}
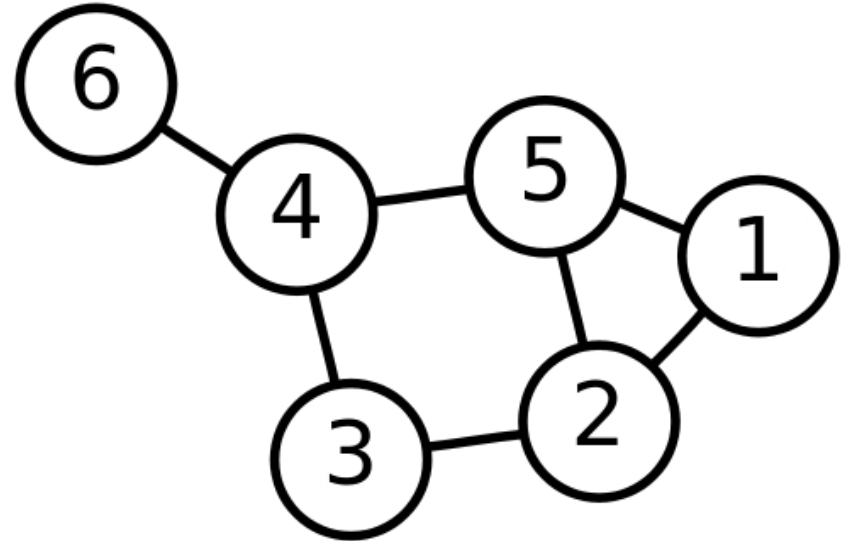
Directed Graphs

- Previously assumed that all edges in the graph were two-way
- Now we have one-way edges:
- Nodes: Same as before
- Edges: (order matters)
  - {(6, 4), (4, 5), (5, 1), (5, 2), (2, 1)}

The Theory Behind PageRank

- The Internet is really a directed graph:
    - Nodes: webpages
    - Edges: links between webpages
- Terms (Suppose we have a page A that links to page B):
    - Page A has a <u>forward-link</u> to page B
    - Page B has a <u>back-link</u> from page A

The Magic Formula

$$R'(u) = c \sum_{v \in B_u} \frac{R'(v)}{N_v} + cE(u)$$

The Magic Formula

Node u is the vertex (webpage) we're interested in computing the PageRank of

$$R'(u) = c \sum_{v \in B_u} \frac{R'(v)}{N_v} + cE(u)$$

The Magic Formula

R'(u) is the PageRank of Node *u*

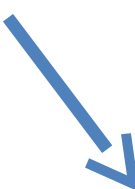$$R'(u) = c \sum_{v \in B_u} \frac{R'(v)}{N_v} + cE(u)$$

The Magic Formula

c is a normalization factor that we can ignore for our purposes

$$R'(u) = c \sum_{v \in B_u} \frac{R'(v)}{N_v} + cE(u)$$

The Magic Formula

E(u) is a "personalization" factor that we can ignore for our purposes

$$R'(u) = c \sum_{v \in B_u} \frac{R'(v)}{N_v} + cE(u)$$

The Magic Formula

We sum over all backlinks of *u*: the PageRank of the website *v* linking to *u* divided by the number of forward-links that *v* has

$$R'(u) = c \sum_{v \in B_u} \frac{R'(v)}{N_v} + cE(u)$$

But wait! This is Recursive!

- Uh oh! We have a recursive formula with no base-case
- We rely on convergence
    - Choose some initial PageRank value for each site
    - Simultaneously compute/update PageRanks
    - When our Delta is small between iterations:
        - Stop, call it "good enough"
- An example of fixed point iteration (similar to proj2, actually)

Sounds Easy. Why MapReduce?

- Assume in the best case that we've crawled and captured the internet as a series of (url, outgoing links) pairs

- We need about 50 iterations of the PageRank algorithm for it to converge

- We quickly see that running it on one machine is not viable

Building a Web Index using PageRank

- Scrape Webpages
- Strip out content, keep only links (input is key = url, value = links on page at url)
  - This step is actually pushed into the MapReduce
- Feed into PageRank Mapreduce
- Sort Documents by PageRank
- Post-process to build the indices that our Google RLP example used

Using MapReduce to Compute PageRank, Step 1

Map:
- Input:
  - key = URL of website
  - val = source of website
- Output for each outgoing link:
  - key = URL of website
  - val = outgoing link url

Reduce:
- Input:
  - key = URL of website
  - values = Iterable of all outgoing links from that website
- Output:
  - key = URL of website
  - value = Starting PageRank, Outgoing links from that website

Using MapReduce to Compute PageRank, Step 2

Map:

- Input:
  - key = URL of website
  - val = PageRank, Outgoing links from that website

- Output for each outgoing link:
  - key = Outgoing Link URL
  - val = Original Website URL, PageRank, # Outgoing links

Reduce:

- Input:
  - key = Outgoing Link URL
  - values = Iterable of all links to Outgoing Link URL

- Output:
  - key = Outgoing Link URL
  - value = Newly computed PageRank (using the formula), Outgoing links from document @ Outgoing Link URL

Repeat this step until PageRank converges – chained MapReduce!

Using MapReduce to Compute PageRank, Step 3

- Finally, we want to sort by PageRank to get a useful index

- MapReduce's built in sorting makes this easy!

Map:
- Input:
  - key = Website URL
  - value = PageRank, Outgoing Links
- Output:
  - key = PageRank
  - value = Website URL

Using MapReduce to Compute PageRank, Step 3

Reduce:
- In case we have duplicate PageRanks
- Input:
  - key = PageRank
  - value = Iterable of URLs with that PageRank
- Output (for each URL in the Iterable):
  - key = PageRank
  - value = Website URL

Since MapReduce automatically sorts the output from the reducer and joins it together We're done!

Using the PageRanked Index

- Do our usual keyword search with RLP implemented

- Take our results, sort by our pre-generated PageRank values

- Send results to user!

- PageRank is still the basis for Google Search
  - (of course, there are many proprietary enhancements in addition)

Question: Which of the following statements is **<u>TRUE</u>** about our implementation of the MapReduce PageRank Algorithm?

(B) We can exactly compute PageRank using our algorithm

(G) Our algorithm chains multiple map/reduce calls to compute our final set of PageRanks

(P) PageRank generally converges after about 10 iterations of our algorithm

(Y) **PageRank was designed to be implemented using the MapReduce framework**

Summary

- MapReduce Data Parallelism
  - Divide large data set into pieces for independent parallel processing
  - Combine and process intermediate results to obtain final result
- Simple to Understand
  - But we can still build complicated software
  - Chaining lets us use the MapReduce paradigm for many common graph and mathematical tasks
- MapReduce is a "Real-World" Tool
  - Worker restart, monitoring to handle failures
  - Google PageRank, Facebook Analytics

Further Reading (Optional)

- Some PageRank slides adapted from http://www.cs.toronto.edu/~jasper/PageRankForMapReduceSmall.pdf
- PageRank Paper:
  - Lawrence Page, Sergey Brin, Rajeev Motwani, Terry Winograd.
    The PageRank Citation Ranking: Bringing Order to the Web.