

Systems Infrastructure for Data Science

Web Science Group

Uni Freiburg

WS 2013/14

Lecture XI:

MapReduce & Hadoop

The new world of Big Data
(programming model)

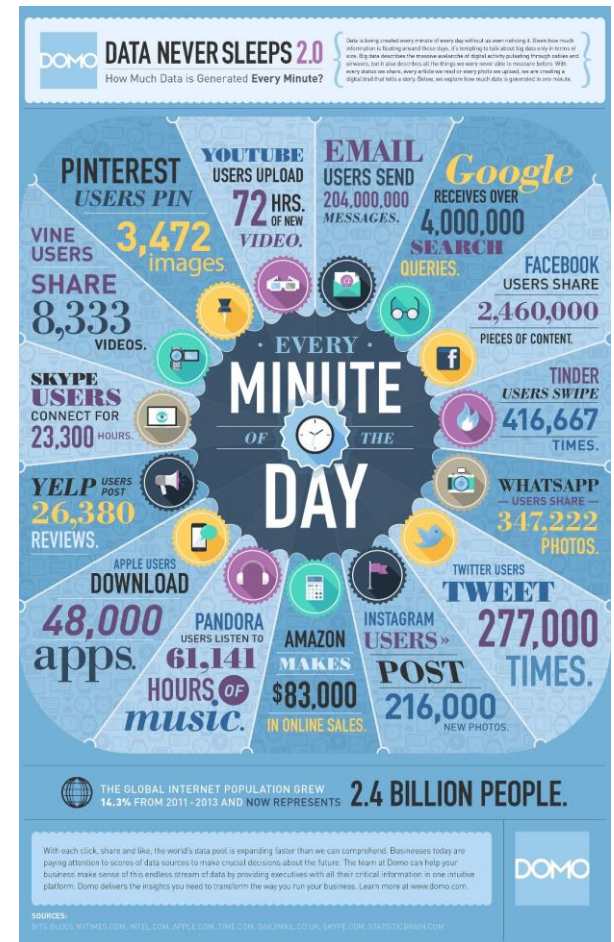
Big Data

- Buzzword for challenges occurring in current data management and analysis
- No longer just storage and retrieval, but also complex computations
- Often expressed as the 4 -7 (depending on source) V's

1st V: Volume

- Scale of Data
 - Scientific applications
(CERN: 70MPixel*40M/s,
15PB/year)
 - Genomics:
(single genome > 1.5TB)
 - Web Data
 - ...

90% of all data was created
in the last two years!



2nd V: Velocity

Speed of data and expected reactions

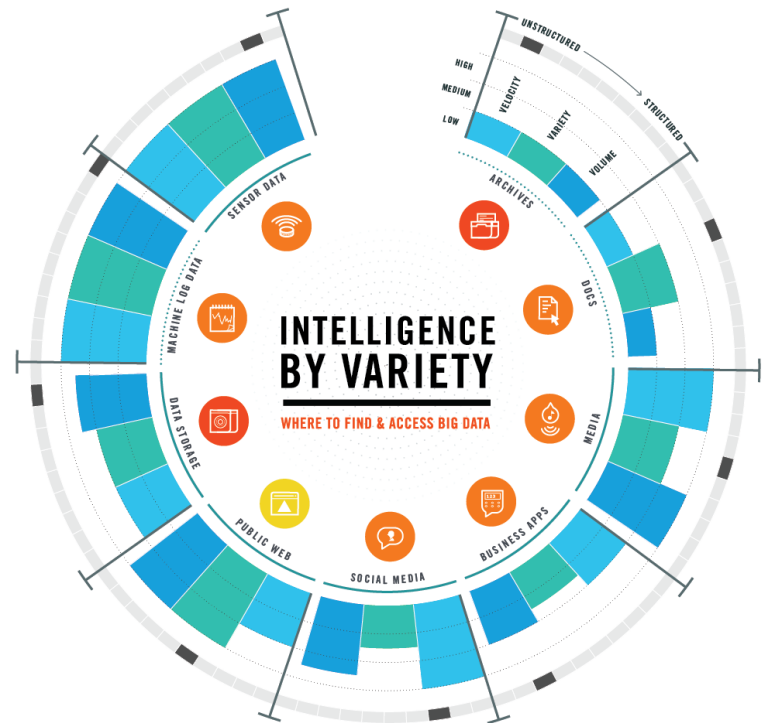
- Stock exchanges (NASDAQ: >35K msg/s, 1ms for common operations)
- Social Media (>150K msg/s peak on Twitter)
- Environmental Sensors (>100 sensors on a car, ms response time)
- Web indexing (reindex within minutes, queries with less than 0.5 seconds)



3rd V: Variety

Form(at) of data not uniform

- Structured vs non-structured (or hidden structure):
relations,
graph,
text,
audio/voice, video, ...
- Broad range of sources:
customers, transactions,
logs, sensors, ...



4th V: Veracity

Uncertainty of Data

- Data Quality and Completeness

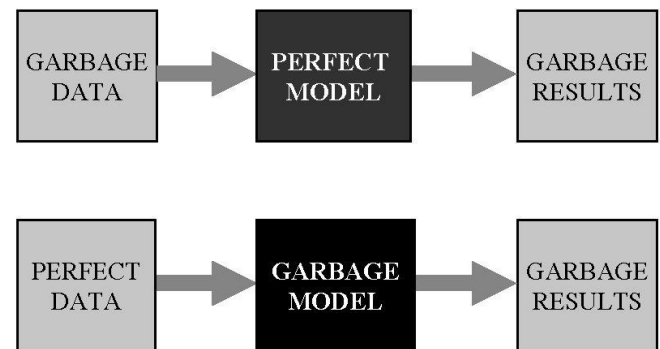
- Sensor readings inconsistent (faults, calibration, ...)
- Social media messages contain slang, abbreviations, colloquialism, ...
- User Profiles faked, duplicated, ...

- Interpretation

- Underlying Model unknown
- Wrong choice of parameters

MODEL CALCULATIONS

“Garbage In-garbage Out” Paradigm



<http://blog.potterzot.com/2007/09/25/garbage-in-garbage-out-and-the-desire-to-cover-our-own-ass-is-ruining-the-world/>

Additional/Disputed V's

- **Value:**
does the analysis yield results that can be used/has applications (customer analysis, trading, business/personal/technology/... improvement)
- **Variability:**
properties of data change over time
- **Visualization:**
complex data cannot be understood without appropriate presentation

Are databases the right tool for these challenges?



<http://dilbert.com/strips/comic/1996-02-27/>

While databases have many benefits, they only serve specific niches

Conceptual Limitations of Relational DBMS

- Well-defined but strict data model
(unordered relations)
- Well-optimizable but limited expression set
(queries+updates)
- Single, strong transaction model
(ACID, Serializability)
- Single interaction model
(store, then query)

Practical Limitations of Traditional Relational DBMS

- Limited Scalability (dozens of nodes)
- Parallel operations low-level or as part of extremely expensive licensing
- Little extensibility (Imperative SQL)
- Labour-intensive to maintain and tune
- Disk-based operations

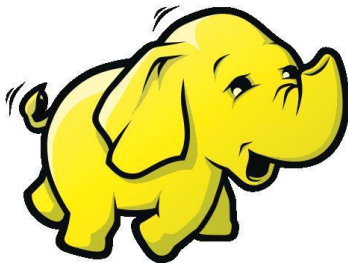
Computational Setting for Big Data

- Computations that need the power of many computers
 - large datasets
 - use of thousands of CPUs in parallel
- Big data management, storage, and analytics
 - cluster as a computer



What is Hadoop?

- Hadoop is an ecosystem of tools for processing “Big Data”.
- Hadoop is an open source project.
- Picks up and advances many design ideas and architectures by big data research
- Starting Point: Google Map/Reduce



Google Cluster Architecture: Key Ideas

- **Single-thread performance doesn't matter**
 - For large problems, **total throughput/\$** is more important than peak performance.
 - **Stuff breaks**
 - If you have 1 server, it may stay up three years (1,000 days).
 - If you have 10,000 servers, expect to lose 10 per day.
 - **“Ultra-reliable” hardware doesn't really help**
 - At large scales, the most reliable hardware still fails, albeit less often
 - Software still needs to be fault-tolerant
 - Commodity machines without fancy hardware give better **performance/\$**
-
- Have a reliable computing infrastructure from clusters of unreliable commodity PCs.
 - Replicate services across many machines to increase request throughput and availability.
 - Favor price/performance over peak performance.

Hadoop: Architectural Design Principles

- Linear scalability
 - More nodes can do more work within the same time
 - Linear on data size, linear on compute resources
- Move computation to data
 - Minimize expensive data transfers
 - Data is large, programs are small
- Reliability and Availability: Failures are common
- Simple computational model (MapReduce)
 - Hides complexity in efficient execution framework
- Streaming data access (avoid random reads)
 - More efficient than seek-based data access

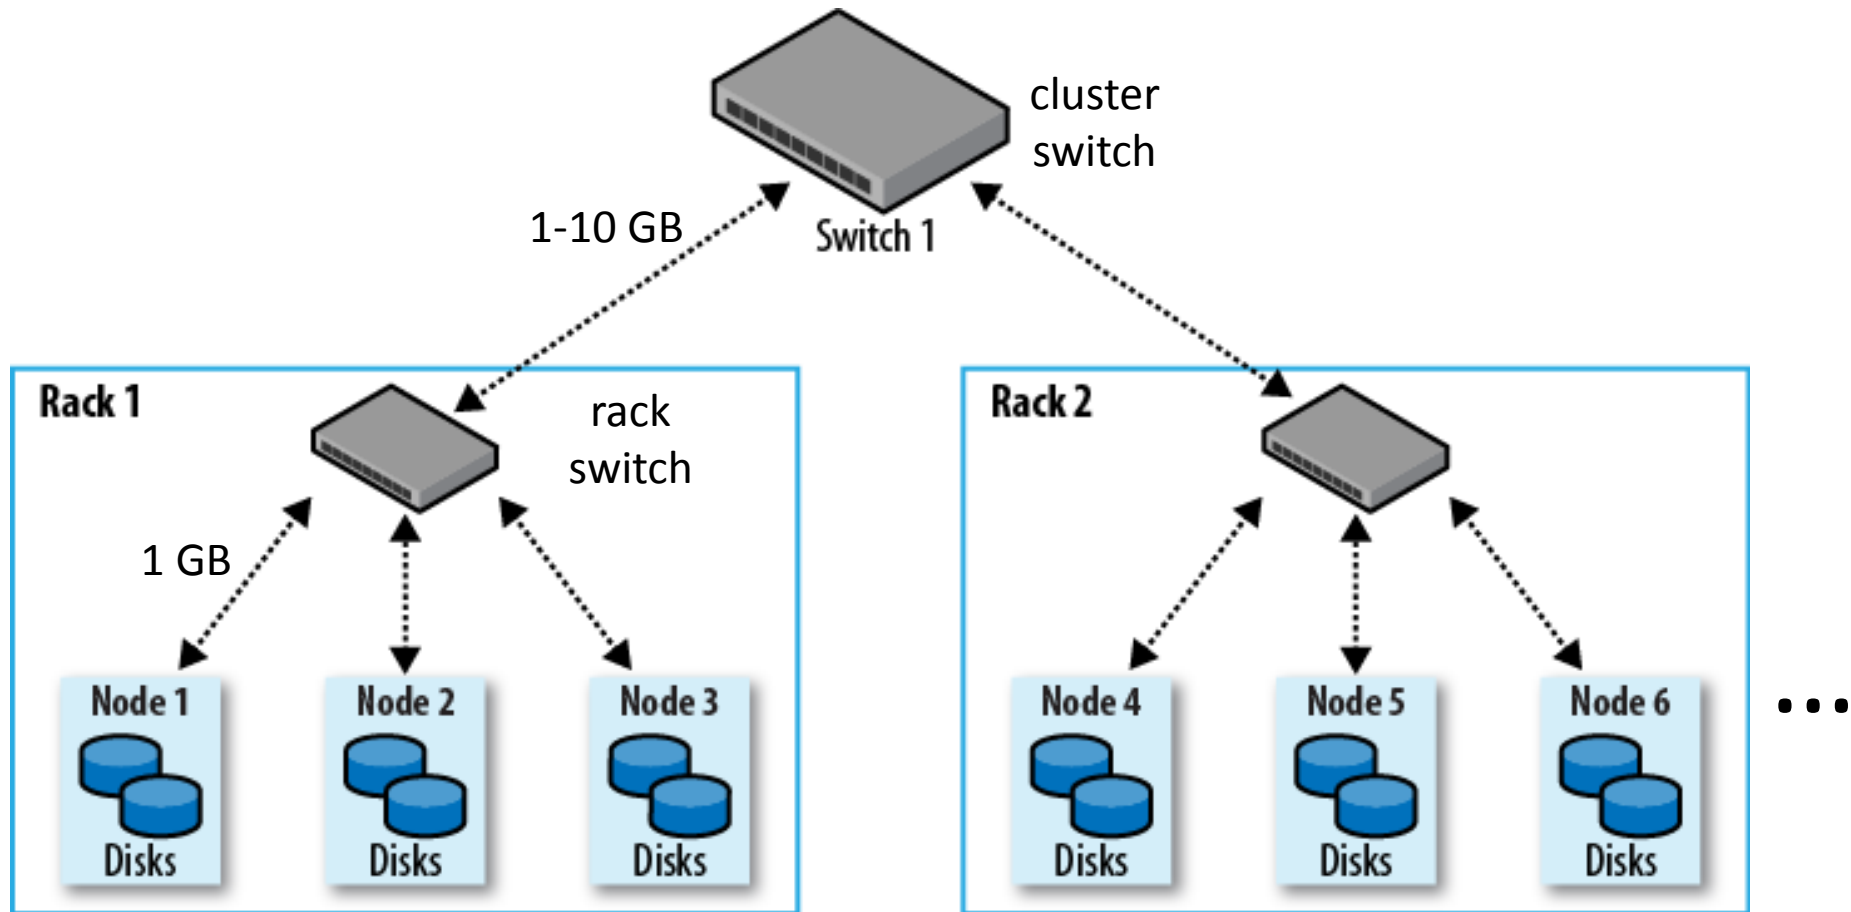
Overview of this Lecture Module

- Background
- Cluster File Storage
 - GFS
 - HDFS
- Computation: MapReduce
 - Model
 - Implementation
 - Programming and Usage
- The Hadoop Ecosystem (next lecture)

The Hadoop Family

MapReduce	Distributed computation framework (data processing model and execution environment)
HDFS	Distributed file system
<i>YARN</i>	<i>Cluster Resource Management</i>
<i>HBase</i>	<i>Distributed, column-oriented database</i>
<i>Hive</i>	<i>Distributed data warehouse</i>
<i>Pig</i>	<i>Higher-level data flow language and parallel execution framework</i>
<i>Mahout</i>	<i>Machine learning and data mining library</i>
ZooKeeper	Distributed coordination service
Avro	Data serialization system (RPC and persistent data storage)
Chukwa	System for collecting management data
BigTop	Packaging and testing

A Typical Hadoop Cluster Setup



~ 30-40 servers per rack

Overview of this Lecture Module

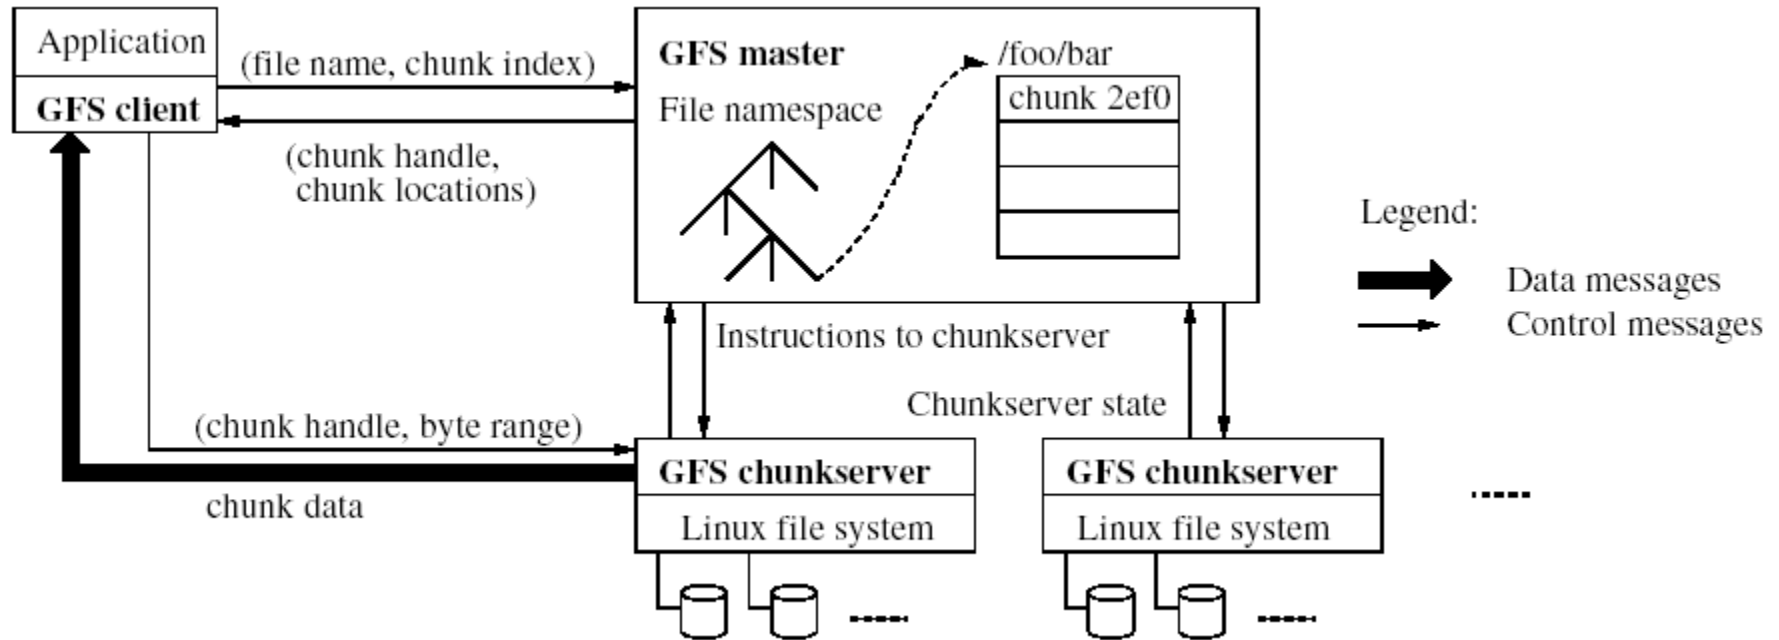
- Background
- Cluster File Storage
 - GFS
 - HDFS
- Computation: MapReduce
 - Model
 - Implementation
 - Programming and Usage
- The Hadoop Ecosystem (next lecture)

Google File System (GFS) Architecture

- Files divided into fixed-sized chunks (64 MB)
 - Each chunk gets a chunk handle from the master
 - Stored as Linux files
- One master
 - Maintains all file system metadata
 - Talks to each chunkserver periodically
- Multiple chunkservers
 - Store chunks on local disks
 - No caching of chunks (not worth it)
- Multiple clients
 - Clients talk to the master for metadata operations
 - Metadata can be cached at the clients
 - Read / write data from chunkservers

GFS Architecture

- Single master, multiple chunkservers



- To overcome single-point of failure & scalability bottleneck:
 - Use shadow masters
 - Minimize master involvement (large chunks; use only for metadata)

Hadoop Distributed File System (HDFS)

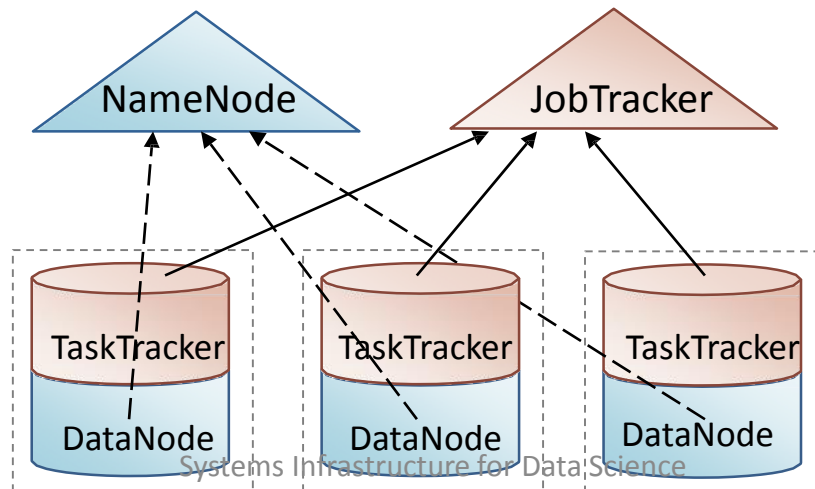
- Hadoop has a general-purpose file system abstraction (i.e., it can use local files, HDFS, Amazon S3, etc.)
- HDFS is Hadoop's file system, implements GFS ideas
- Streaming data access
 - write-once, read-many-times pattern
 - time to read the whole dataset is more important
- HDFS is not a good fit for
 - low-latency data access
 - lots of small files
 - multiple writers, arbitrary file modifications

Blocks

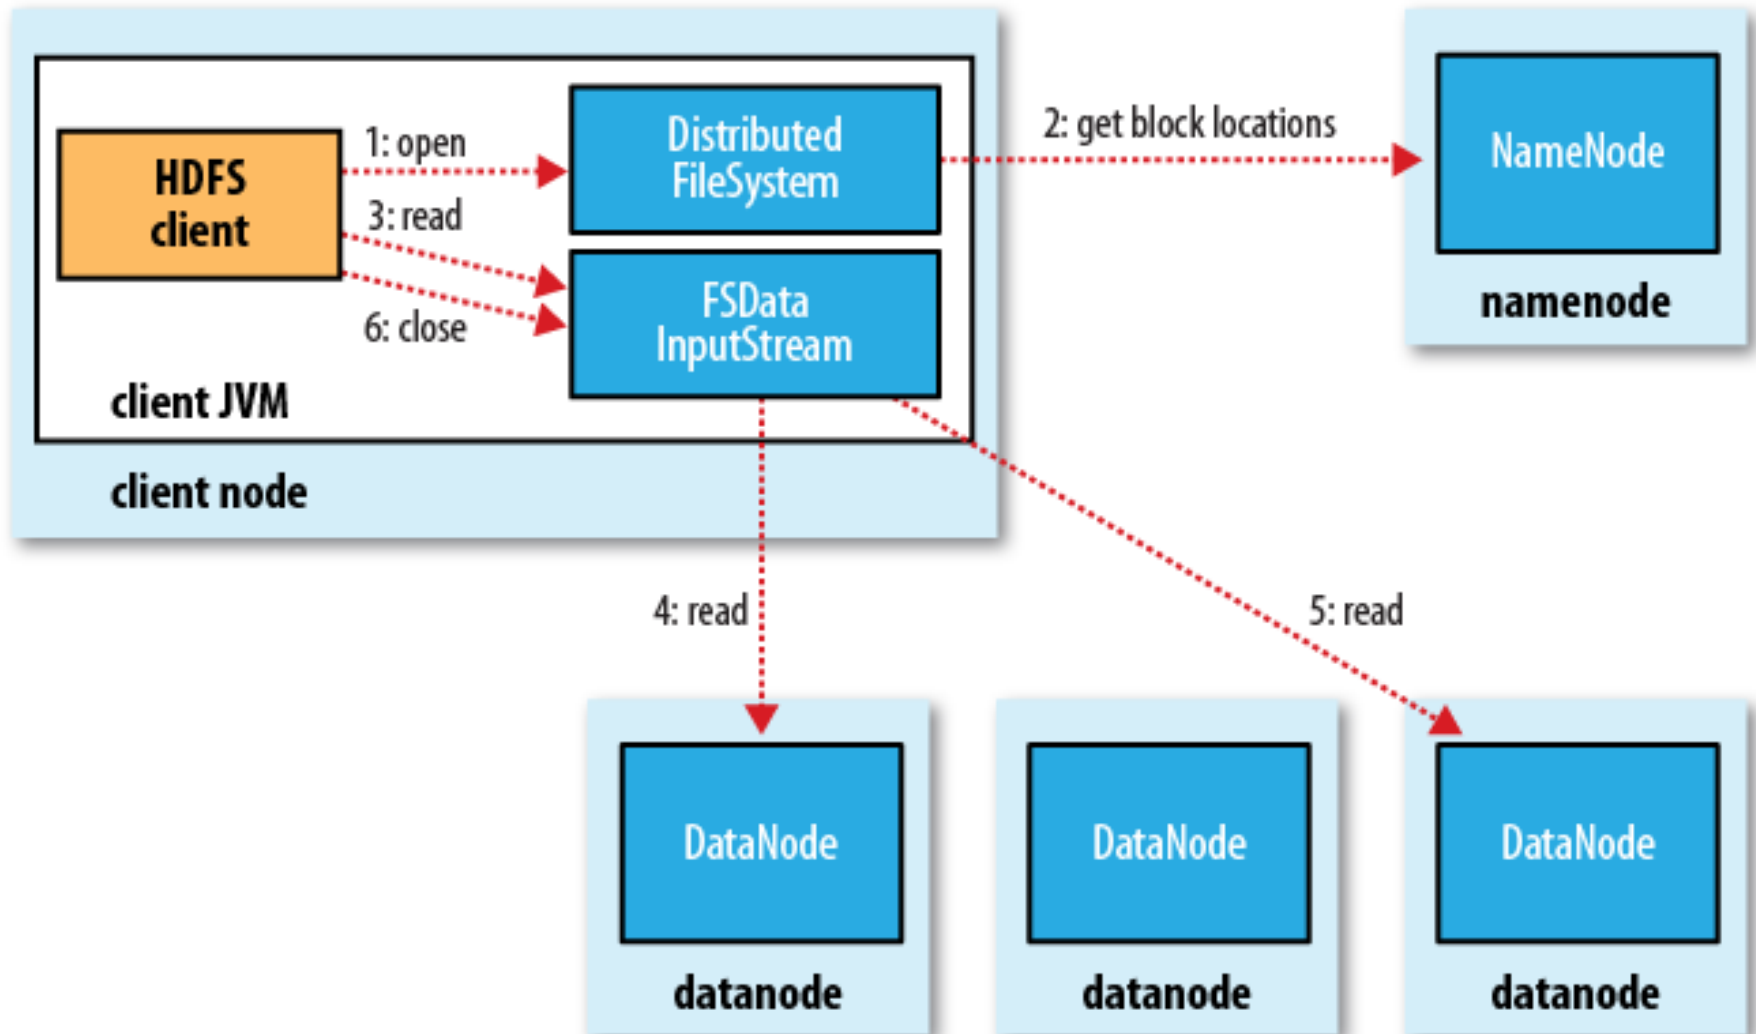
- HDFS files are broken into block-sized chunks (64 MB by default)
- With the (large) block abstraction:
 - a file can be larger than any single disk in the network
 - storage subsystem is simplified (e.g., metadata bookkeeping)
 - replication for fault-tolerance and availability is facilitated

Hadoop Main Cluster Components

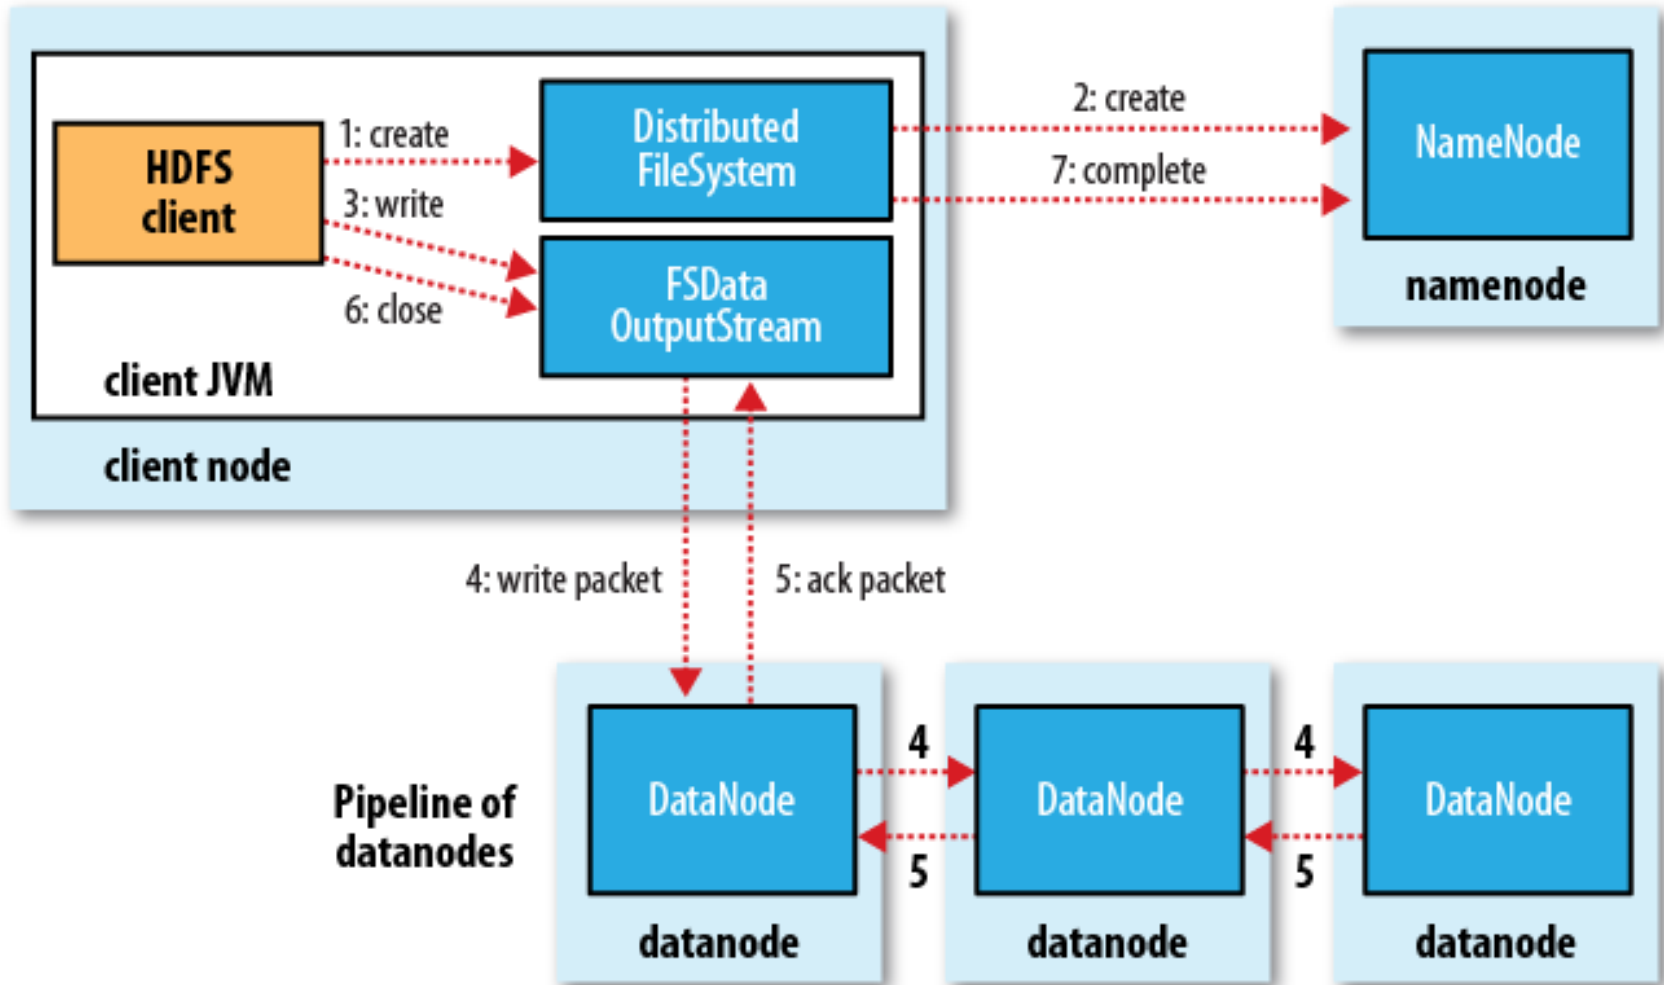
- HDFS daemons
 - **NameNode**: namespace and block management (~ master in GFS)
 - **DataNodes**: block replica container (~ chunkserver in GFS)
- MapReduce daemons
 - **JobTracker**: client communication, job scheduling, resource management, lifecycle coordination (~ master in Google MR)
 - **TaskTrackers**: task execution module (~ worker in Google MR)



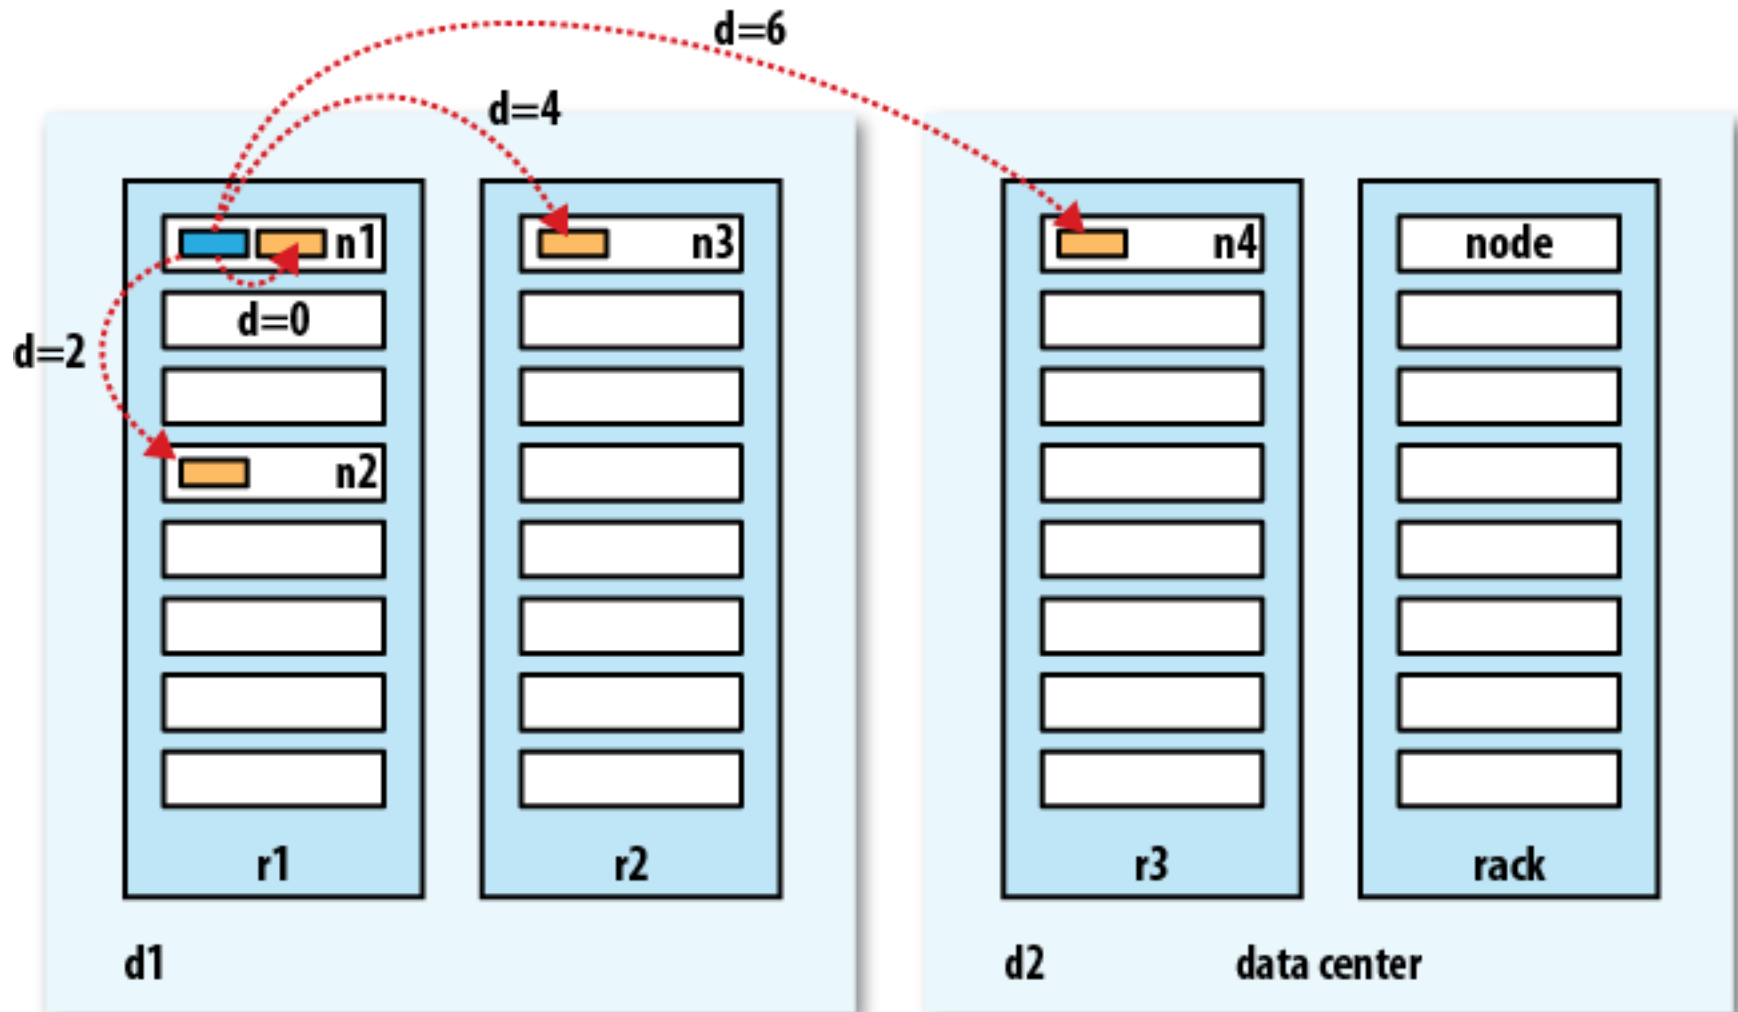
Reading from HDFS



Writing to HDFS



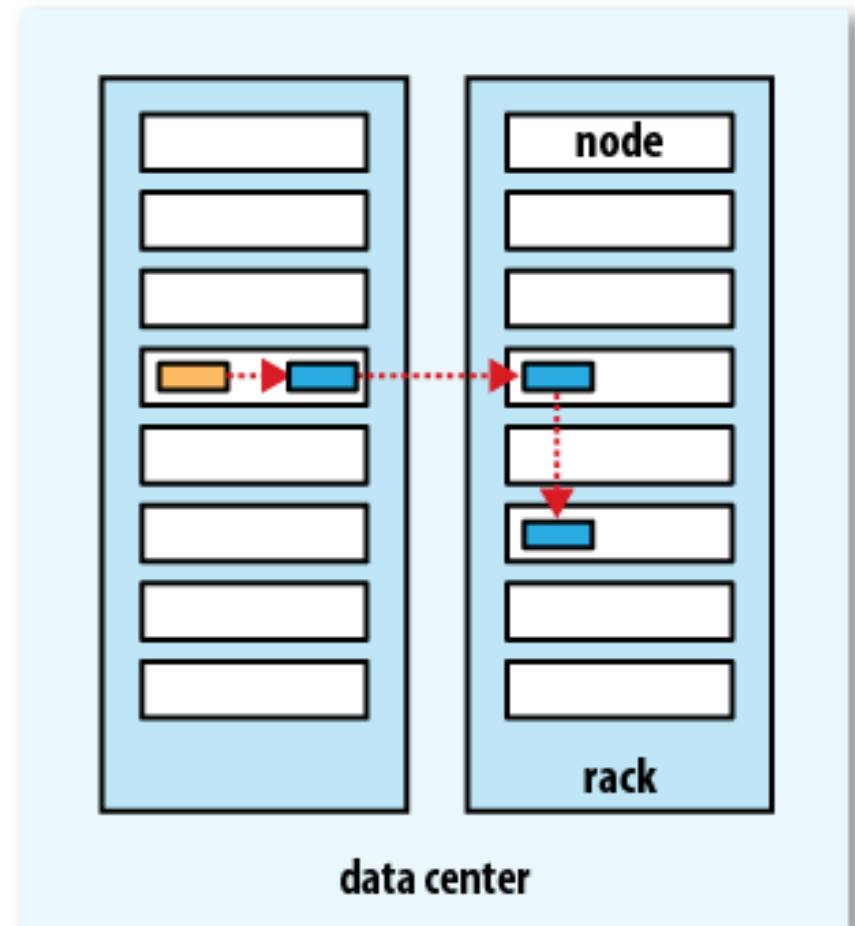
Network Distance in Hadoop



Replica Placement

- Issues to consider: reliability, write bandwidth, read bandwidth, block distribution.
- Hadoop's default strategy:
 - First replica: on the client node (or randomly chosen if client is outside the cluster)
 - Second replica: random, off-rack.
 - Third replica: same rack as second, different node.
 - More replicas: randomly chosen.

Example Replica Pipeline:



Coherency Model

- Coherency model describes the data visibility of reads and writes for a file.
- In HDFS:
 - The metadata for a newly created file is visible in the file system namespace.
 - The current data block being written is not guaranteed to be visible to other readers.
- To force all buffers to be synchronized to all relevant datanodes, you can use the `sync()` method.
- Without `sync()`, you may lose up to a block of (newly written) data in the event of client or system failure.

HDFS Federation & High-Availability

- In latest releases of Hadoop:
 - **HDFS Federation** allows multiple Namenodes, each of which manages a portion of the file system namespace; the goal is to enhance the scalability of the Namenode on very large clusters with many files and blocks.
 - **HDFS High-Availability** provides faster recovery from Namenode failures using a pair of namenodes in an active standby configuration.

Overview of this Lecture Module

- Background
- Cluster File Storage
 - GFS
 - HDFS
- Computation: MapReduce
 - Model
 - Implementation
 - Programming and Usage
- The Hadoop Ecosystem (next lecture)

MapReduce

- a **software framework** first introduced by Google in 2004 to support parallel and fault-tolerant computations over large data sets on clusters of computers
- based on the **map/reduce functions** commonly used in the functional programming world

MapReduce in a Nutshell

- Given:
 - a very large dataset
 - a well-defined computation task to be performed on elements of this dataset (preferably, in a parallel fashion on a large cluster)
- MapReduce framework:
 - Just express what you want to compute (`map()` & `reduce()`).
 - Don't worry about parallelization, fault tolerance, data distribution, load balancing (MapReduce takes care of these).
 - What changes from one application to another is the actual computation; the programming structure stays similar.

MapReduce in a Nutshell

- Here is the framework in simple terms:
 - Read lots of data.
 - **Map**: extract something that you care about from each record.
 - Shuffle and sort.
 - **Reduce**: aggregate, summarize, filter, or transform.
 - Write the results.
- One can use as many Maps and Reduces as needed to model a given problem.

MapReduce vs. Traditional RDBMS

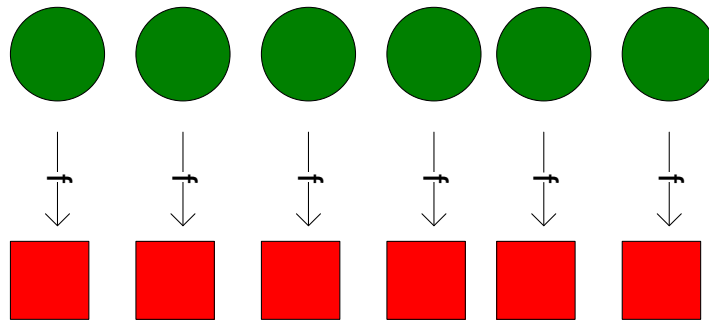
	MapReduce	Traditional RDBMS
Data size	Petabytes	Gigabytes
Access	Batch	Interactive and batch
Updates	Write once, read many times	Read and write many times
Structure	Dynamic schema	Static schema
Integrity	Low	High (normalized data)
Scaling	Linear	Non-linear (general SQL)

Functional Programming Foundations

- map in MapReduce \leftrightarrow map in FP
- reduce in MapReduce \leftrightarrow fold in FP
- Note: There is no precise 1-1 correspondence, but the general idea is similar.

map() in Haskell

- Create a new list by applying f to each element of the input list.



- **Definition of map:**

`map :: (a → b) → [a] → [b]` -- type of map

`map f [] = []` -- the empty list case

`map f (x:xs) = f x : map f xs` -- the non-empty list case

- **Example: Double all numbers in a list.**

Haskell-prompt `> map ((* 2) [1, 2, 3]`

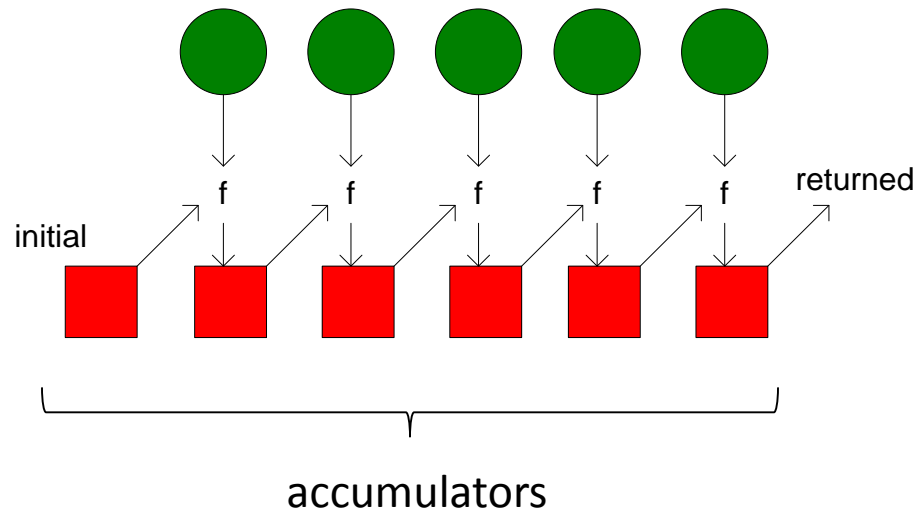
`[2, 4, 6]`

Implicit Parallelism in map()

- In a purely functional setting, an element of a list being computed by map cannot see the effects of the computations on other elements.
- If the order of application of a function f to elements in a list is commutative, then we can reorder or parallelize execution.
- This is the “secret” that MapReduce exploits.

fold() in Haskell

- Move across a list, applying a function **f** to each element plus an **accumulator**. **f** returns the next accumulator value, which is combined with the next element of the list.



- Two versions: fold left & fold right

fold() in Haskell

- **Definition of fold left:**

`foldl :: (b → a → b) → b → [a] → b` -- type of foldl

`foldl f y [] = y` -- the empty list case

`foldl f y (x:xs) = foldl f (f y x) xs` -- the non-empty list case

- **Definition of fold right:**

`foldr :: (a → b → b) → b → [a] → b` -- type of foldr

`foldr f y [] = y` -- the empty list case

`foldr f y (x:xs) = f x (foldr f y xs)` -- the non-empty list case

- **Example: Compute the sum of all numbers in a list.**

Haskell-prompt > foldl (+) 0 [1, 2, 3]	}	<code>foldl (+) 0 [1, 2, 3]</code>
6		<code>⇒ (((0 + 1) + 2) + 3)</code>
		<code>⇒ 6</code>

reduce() in Haskell

- reduce is a type-specialized version of fold.

- **Definition of reduce:**

`reduce :: (a → a → a) → a → [a] → a` -- type of reduce

`reduce = foldl` -- definition of reduce

MapReduce Basic Programming Model

- Transform a set of input key-value pairs to a set of output values:
 - Map: $(k1, v1) \rightarrow \text{list}(k2, v2)$
 - MapReduce library groups all intermediate pairs with same key together.
 - Reduce: $(k2, \text{list}(v2)) \rightarrow \text{list}(v2)$

MapReduce Canonical Example

“Count word occurrences in a set of documents.”

map(k1, v1) → list(k2, v2)

map (String key, String value):
// key: document name
// value: document contents
for each word w in value:
 EmitIntermediate(w, "1");

“document1”, “to be or not to be”

↓
“to”, “1”
“be”, “1”
“or”, “1”
...

reduce(k2, list(v2)) → list(v2)

reduce (String key, Iterator values):
// key: a word
// values: a list of counts
int result = 0;
for each v in values:
 result += ParseInt(v);
Emit(AsString(result));

key = “be”
values = “1”, “1”

↓
“2”

key = “not”
values = “1”

↓
“1”

key = “or”
values = “1”

↓
“1”

key = “to”
values = “1”, “1”

↓
“2”

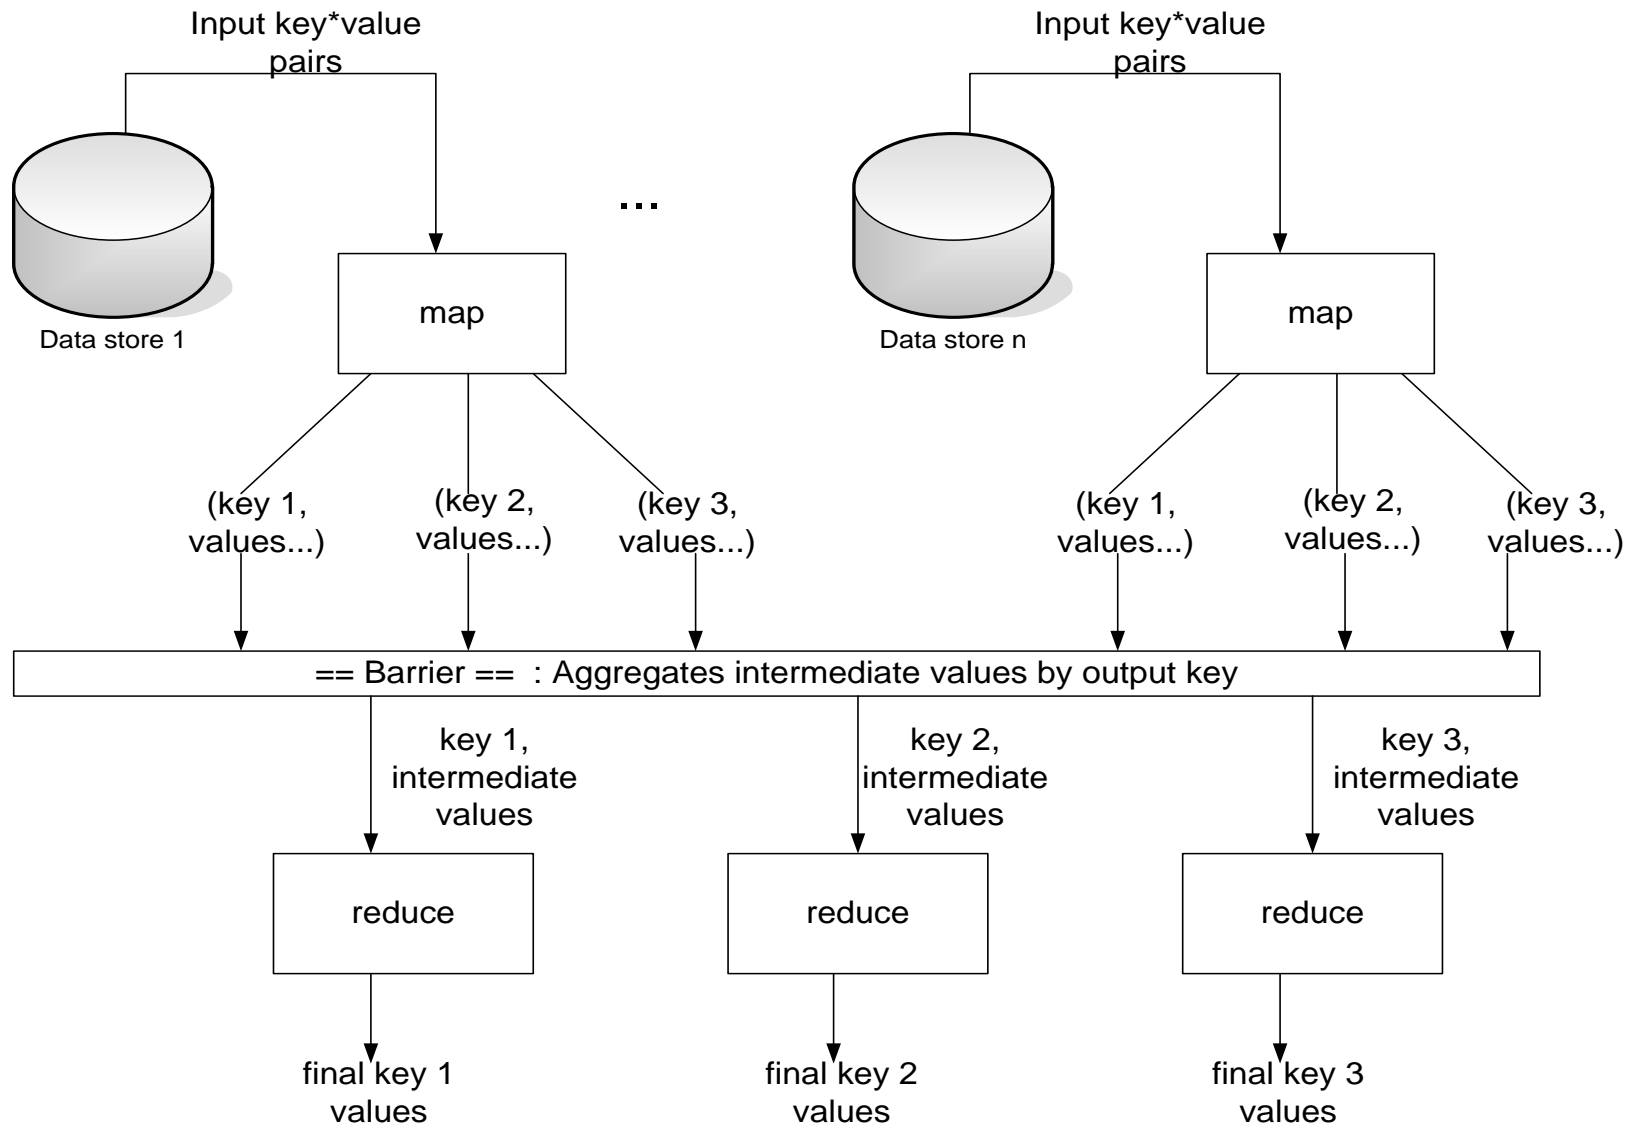
Overview of this Lecture Module

- Background
- Cluster File Storage
 - GFS
 - HDFS
- Computation: MapReduce
 - Model
 - Implementation
 - Programming and Usage
- The Hadoop Ecosystem (next lecture)

MapReduce Parallelization

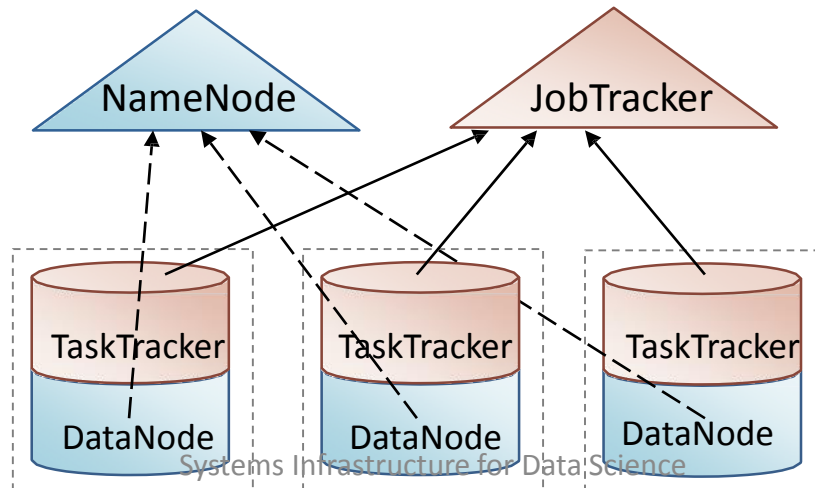
- Multiple map() functions run in parallel, creating different intermediate values from different input data sets.
- Multiple reduce() functions also run in parallel, each working on a different output key.
- All values are processed independently.
- **Bottleneck: The reduce phase can't start until the map phase is completely finished.**

MapReduce Parallel Processing Model

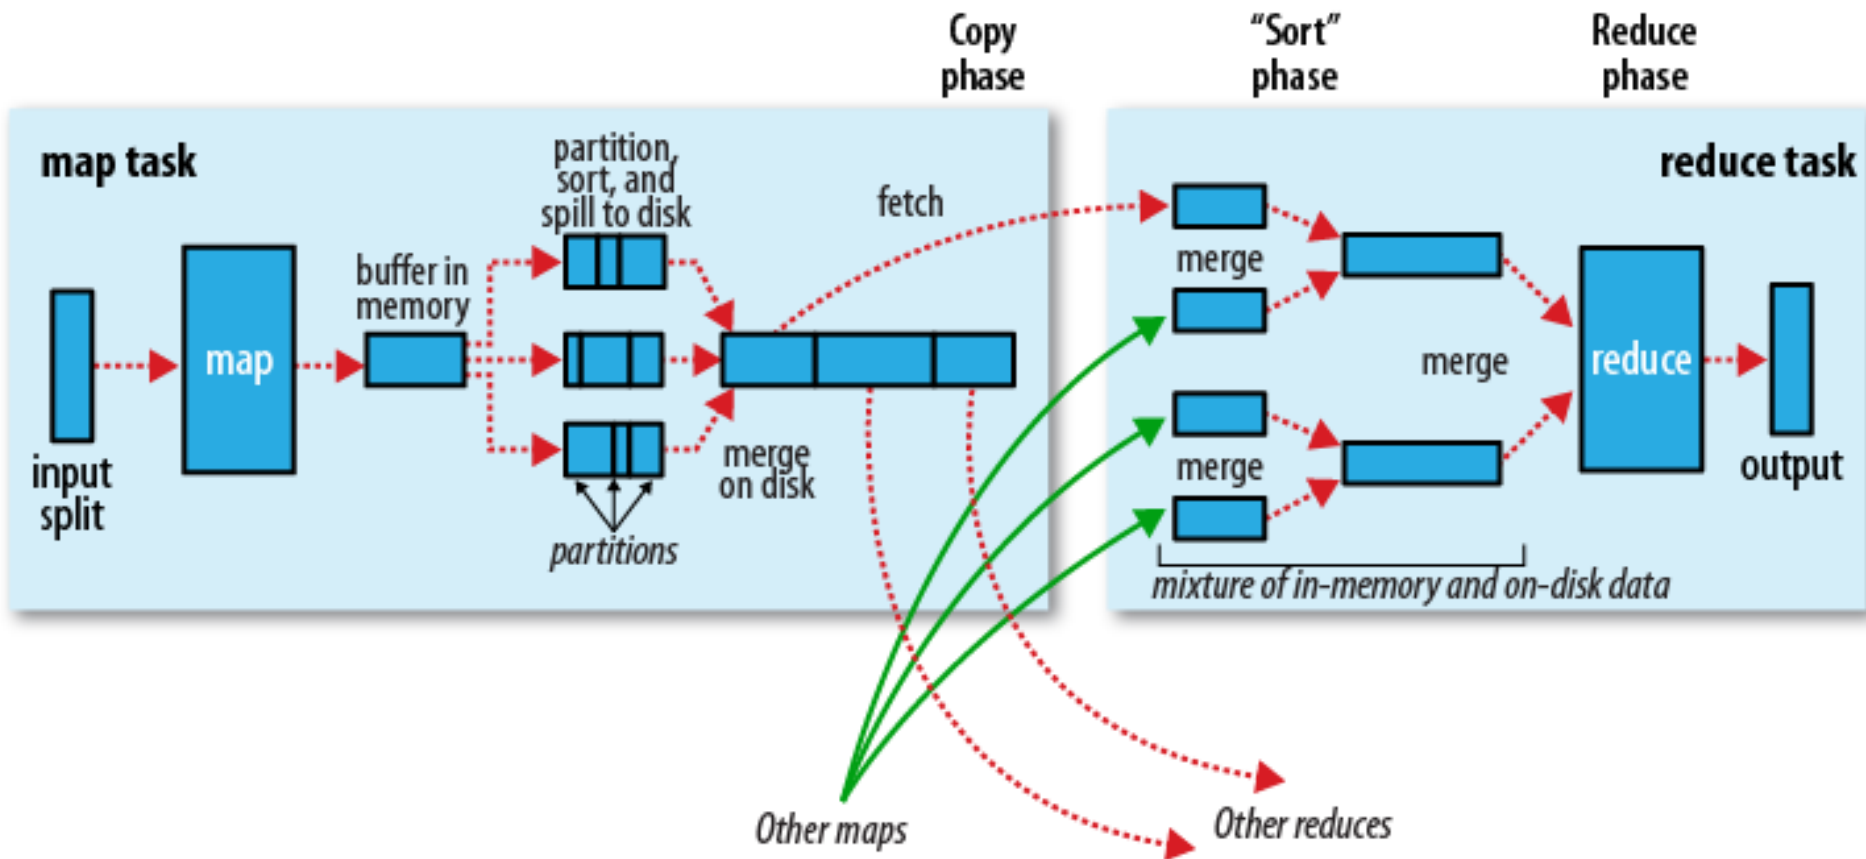


Hadoop Main Cluster Components

- HDFS daemons
 - **NameNode**: namespace and block management (~ master in GFS)
 - **DataNodes**: block replica container (~ chunkserver in GFS)
- MapReduce daemons
 - **JobTracker**: client communication, job scheduling, resource management, lifecycle coordination (~ master in Google MR)
 - **TaskTrackers**: task execution module (~ worker in Google MR)



Shuffle & Sort Overview



Moving Data from Mappers to Reducers

- “Shuffle & Sort” phase
 - synchronization barrier between map and reduce phase
 - one of the most expensive parts of a MapReduce execution
- Mappers need to separate output intended for different reducers
- Reducers need to collect their data from all mappers and group it by key
 - keys at each reducer are processed in order

MapReduce Data Locality

- Goal: To conserve network bandwidth.
- In GFS, data files are divided into 64 MB blocks and 3 copies of each are stored on different machines.
- Master program schedules `map()` tasks based on the location of these replicas:
 - Put `map()` tasks physically on the same machine as one of the input replicas (or, at least on the same rack / network switch).
- This way, thousands of machines can read input at local disk speed. Otherwise, rack switches would limit read rate.

MapReduce Scheduling

- One master, many workers
 - Input data split into M map tasks (typically 64 MB (\sim chunk size in GFS))
 - Reduce phase partitioned into R reduce tasks ($\text{hash}(k) \bmod R$)
 - Tasks are assigned to workers dynamically
- Master assigns each map task to a free worker
 - Considers locality of data to worker when assigning a task
 - Worker reads task input (often from local disk)
 - Worker produces R local files containing intermediate k/v pairs
- Master assigns each reduce task to a free worker
 - Worker reads intermediate k/v pairs from map workers
 - Worker sorts & applies user's reduce operation to produce the output

Choosing M and R

- M = number of map tasks, R = number of reduce tasks
- Larger M, R: creates smaller tasks, enabling easier load balancing and faster recovery (many small tasks from failed machine)
- Limitation: $O(M+R)$ scheduling decisions and $O(M \cdot R)$ in-memory state at master
 - Very small tasks not worth the startup cost
- Recommendation:
 - Choose M so that split size is approximately 64 MB
 - Choose R a small multiple of the number of workers; alternatively choose R a little smaller than #workers to finish reduce phase in one “wave”

MapReduce Fault Tolerance (I)

On worker/ JobTracker failure:

- Master/TaskTracker detects failure via periodic heartbeats.
- Both completed and in-progress map tasks on that worker should be re-executed (→ output stored on local disk).
- Only in-progress reduce tasks on that worker should be re-executed (→ output stored in global file system).
- All reduce workers will be notified about any map re-executions.

MapReduce Fault Tolerance (II)

- On master/JobTracker failure:
 - Google:
State is check-pointed to GFS: new master recovers & continues.
 - Hadoop cannot deal with JobTracker failure
 - Could use Google's proposed JobTracker take-over idea, using ZooKeeper to make sure there is at most one JobTracker
 - Improvements in progress in newer releases...
- Robustness:
 - Example: Lost 1600 of 1800 machines once, but finished fine.

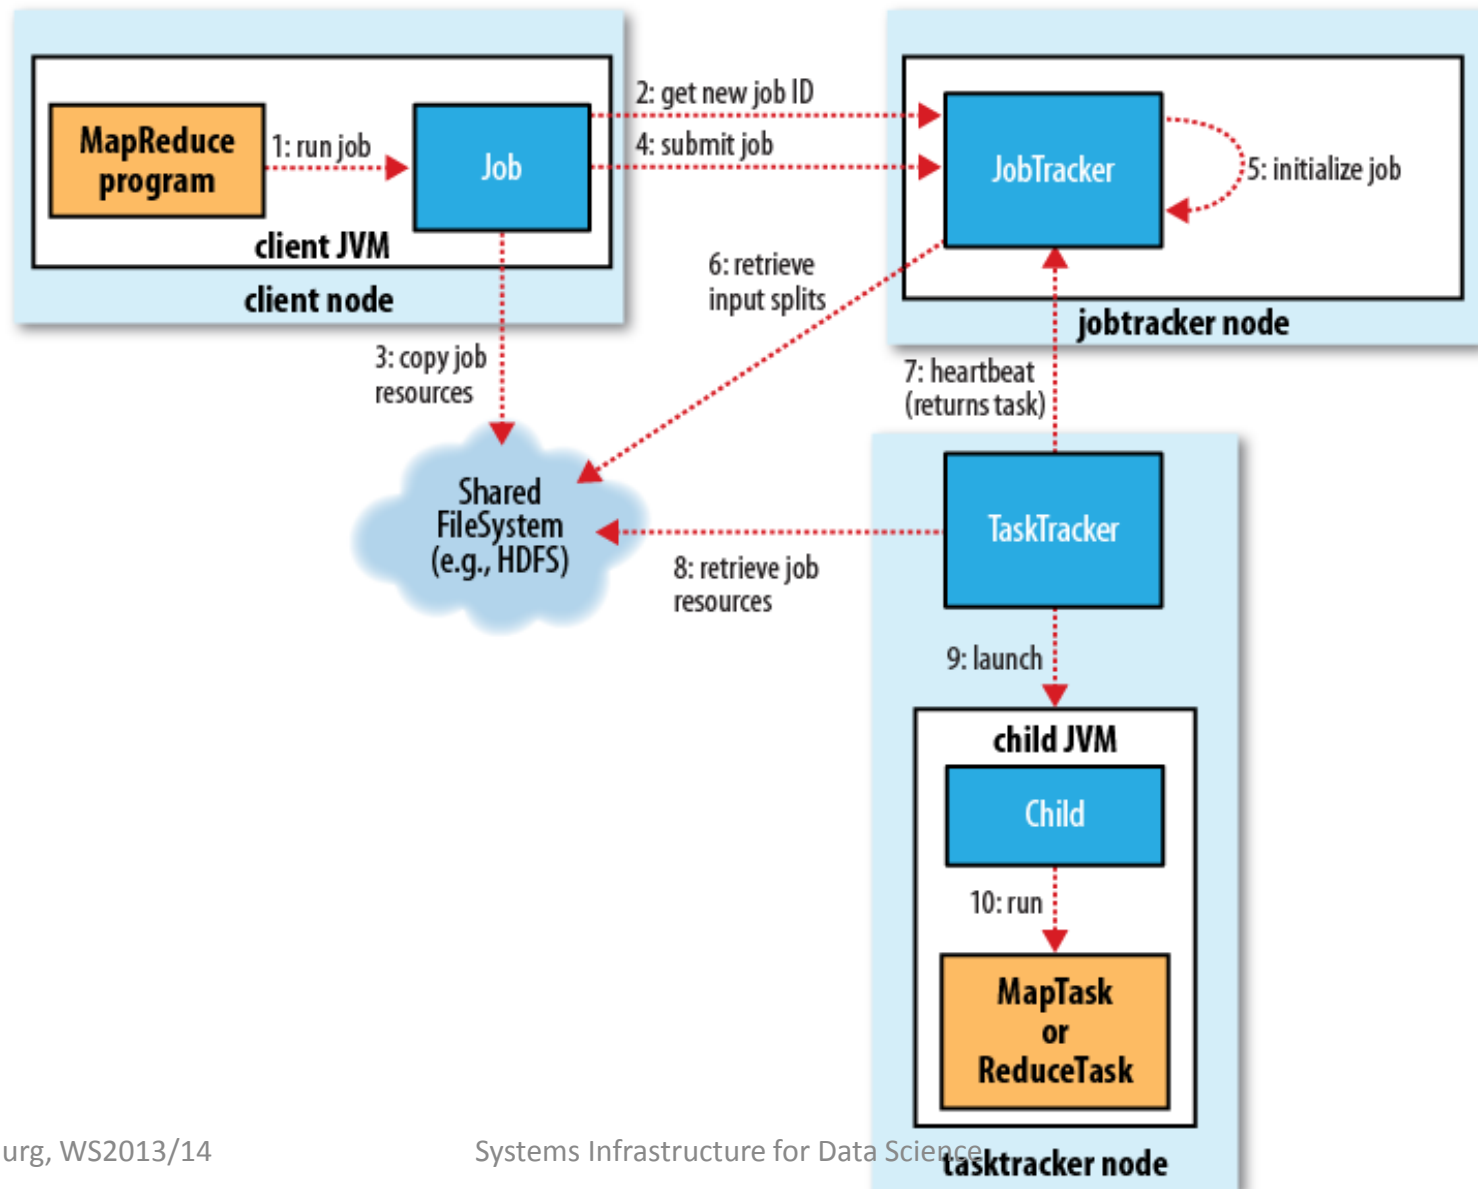
Stragglers & Backup Tasks

- Problem: “Stragglers” (i.e., slow workers) significantly lengthen the completion time.
- Solution: Close to completion, spawn backup copies of the remaining in-progress tasks.
 - Whichever one finishes first, “wins”.
- Additional cost: a few percent more resource usage.
- Example: A sort program without backup = 44% longer.

Overview of this Lecture Module

- Background
- Cluster File Storage
 - GFS
 - HDFS
- Computation: MapReduce
 - Model
 - Implementation
 - Programming and Usage
- The Hadoop Ecosystem (next lecture)

MapReduce Job Execution in Hadoop



Basic MapReduce Program Design

- Tasks that can be performed independently on a data object, large number of them: Map
- Tasks that require combining of multiple data objects: Reduce
- Sometimes it is easier to start program design with Map, sometimes with Reduce
- Select keys and values such that the right objects end up together in the same Reduce invocation
- Might have to partition a complex task into multiple MapReduce sub-tasks

MapReduce Development Steps

- Write Map and Reduce functions
 - Create unit tests
- Write driver program to run a job
 - Can run from IDE with small data subset for testing
 - If test fails, use IDE for debugging
 - Update unit tests and Map/Reduce if necessary
- Once program works on small test set, run it on full data set
 - If there are problems, update tests and code accordingly
- Fine-tune code, do some profiling

Other Practical Extensions

- User-specified **combiner functions** for partial combination within a map task can save network bandwidth (~ mini-reduce)
 - Example: Word Count?
- User-specified **partitioning functions** for mapping intermediate key values to reduce workers (by default: $\text{hash}(\text{key}) \bmod R$)
 - Example: $\text{hash}(\text{Hostname}(\text{urlkey})) \bmod R$
- **Ordering guarantees:** Processing intermediate k/v pairs in increasing order
 - Example: reduce of Word Count outputs ordered results.
- Custom input and output format handlers
- Single-machine execution option for testing & debugging

Combiner Functions

- Pre-reduces mapper output before transfer to reducers (to minimize data transferred)
- Does not change program semantics
- Usually same as reduce function, but has to have same output type as Map
- Works only for certain types of reduce functions (commutative and associative (a.k.a. distributive))
 - E.g.: $\max(5, 4, 1, 2) = \max(\max(5, 1), \max(4, 2))$

Partitioner Functions

- Partitioner determines which keys are assigned to which reduce task
- Default HashPartitioner essentially assigns keys randomly
- Create custom partitioner by implementing your own `getPartition()` method of Partitioner in `org.apache.hadoop.mapreduce`

Local (Standalone) Mode

- Runs same MapReduce user program as cluster version, but does it sequentially on a single machine
- Does not use any of the Hadoop daemons
- Works directly with local file system
 - No HDFS, hence no need to copy data to/from HDFS
- Great for development, testing, initial debugging

Pseudo-Distributed Mode

- Still runs on a single machine, but simulating a real Hadoop cluster
 - Simulates multiple nodes
 - Runs all daemons
 - Uses HDFS
- For more advanced testing and debugging
- You can also set this up on your laptop

Programming Language Support

- Java API (native)
- Hadoop Streaming API
 - allows writing map and reduce functions in any programming language that can read from standard input and write to standard output
 - Examples: Ruby, Python
- Hadoop Pipes API
 - allows map and reduce functions written in C++ using sockets to communicate with Hadoop's TaskTrackers

Overview of this Lecture Module

- Background
- Cluster File Storage
 - GFS
 - HDFS
- Computation: MapReduce
 - Model
 - Implementation
 - Programming and Usage
- The Hadoop Ecosystem (next lecture)

References

- **“Web Search for a Planet: The Google Cluster Architecture”**, L. Barroso, J. Dean, U. Hoelzle, IEEE Micro 23(2), 2003.
- **“The Google File System”**, S. Ghemawat, H. Gobioff, S. Leung, SOSP 2003.
- **“MapReduce: Simplified Data Processing on Large Clusters”**, J. Dean, S. Ghemawat, OSDI 2004 (follow-up papers: CACM 2008, CACM 2010).
- **“The Hadoop Distributed File System”**, K. Shvachko et al, MSST 2010.
- **“Hadoop: The Definitive Guide”**, T. White, O’Reilly, 3rd edition, 2012.