



7610: Distributed Systems

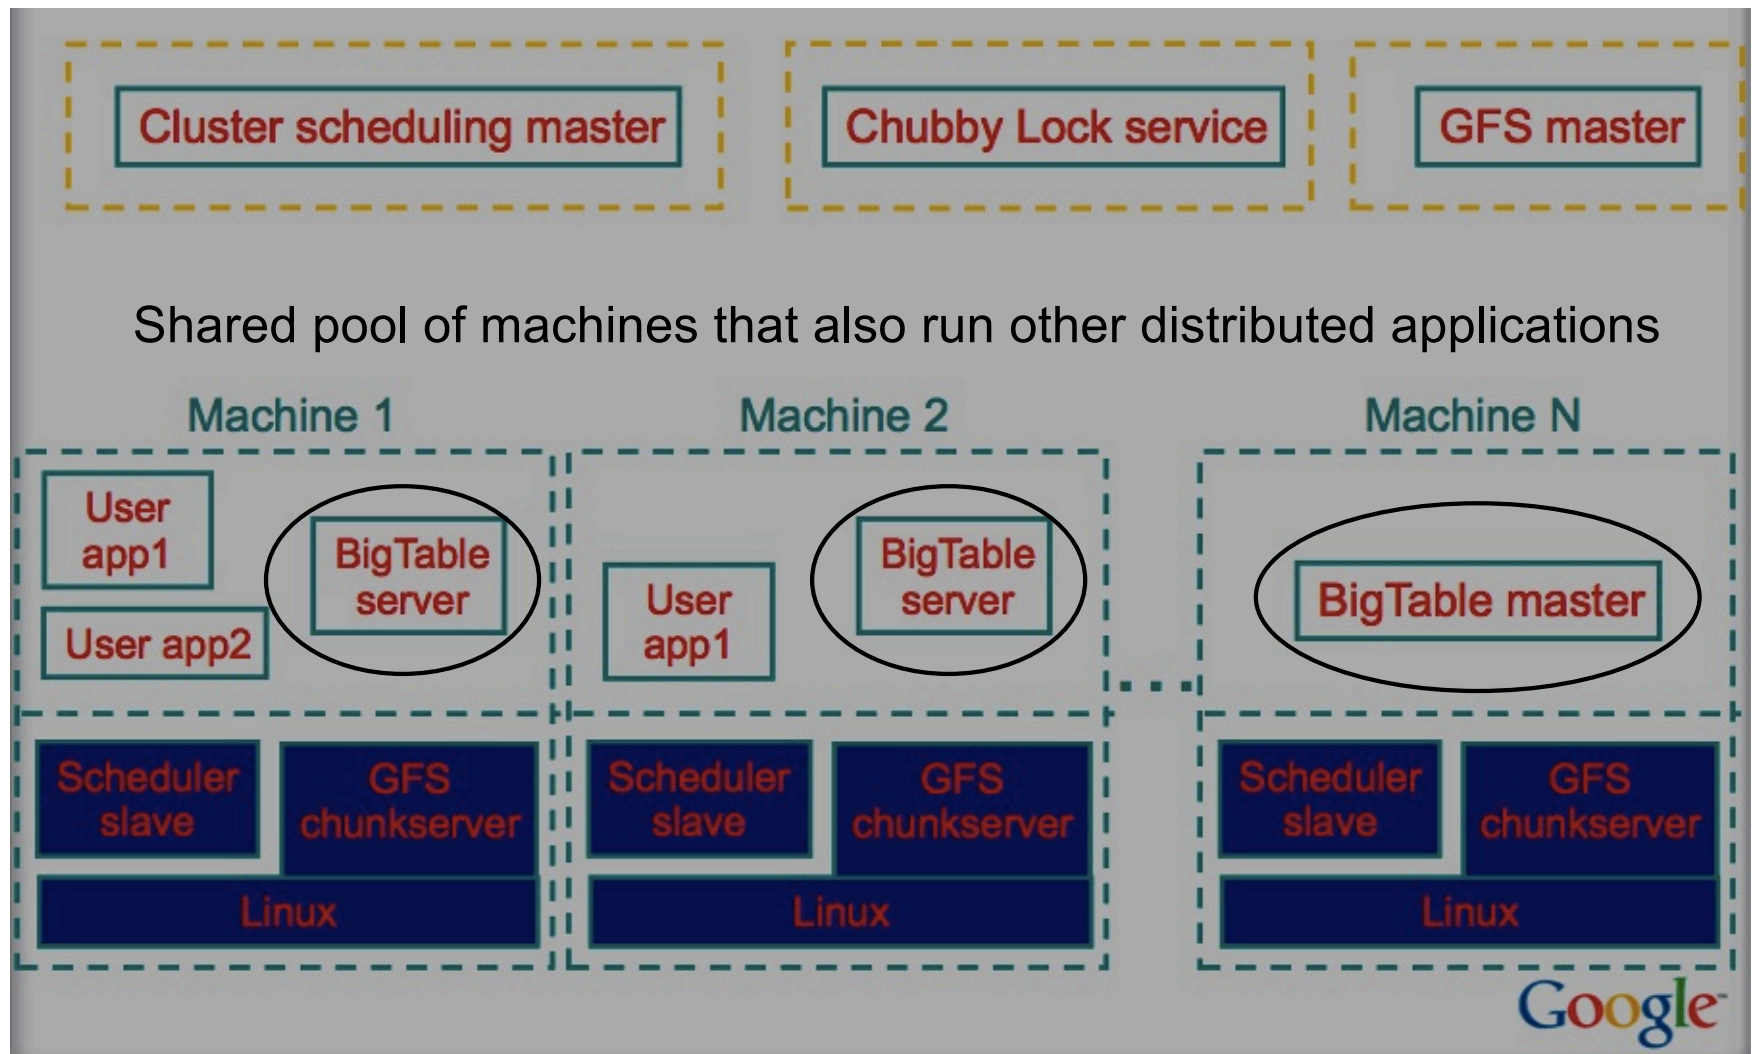
MapReduce. Hadoop. Spark. Mesos. Yarn

REQUIRED READING

- ▶ MapReduce: Simplified Data Processing on Large Clusters OSDI 2004
- ▶ Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center, NSDI 2011
- ▶ Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing, NSDI 2012, (best paper)
- ▶ Apache Hadoop YARN: Yet Another Resource Negotiator SOCC 2013 (best paper)
- ▶ Omega: flexible, scalable schedulers for large compute clusters, EuroSys 2013 (best paper)



Typical Google Cluster





1: MapReduce

These are slides from Dan Weld's class at U. Washington (who in turn made his slides based on those by Jeff Dean, Sanjay Ghemawat, Google, Inc.)

Motivation

- ▶ **Large-Scale Data Processing**
 - ▶ Want to use 1000s of CPUs
 - ▶ But don't want hassle of managing things
- ▶ **MapReduce provides**
 - ▶ Automatic parallelization & distribution
 - ▶ Fault tolerance
 - ▶ I/O scheduling
 - ▶ Monitoring & status updates

Map/Reduce

- ▶ **Map/Reduce**
 - ▶ Programming model from Lisp
 - ▶ (and other functional languages)
- ▶ Many problems can be phrased this way
- ▶ Easy to distribute across nodes
- ▶ Nice retry/failure semantics

Map in Lisp (Scheme)

▶ (map f list [list2 list3 ...])

Unary operator

▶ (map square '(1 2 3 4))

▶ (1 4 9 16)

Binary operator

▶ (reduce + '(1 4 9 16))

▶ (+ 16 (+ 9 (+ 4 1)))

▶ (reduce + (map square (map - 11 12))))

Map/Reduce ala Google

- ▶ `map(key, val)` is run on each item in set
 - ▶ emits new-key / new-val pairs
- ▶ `reduce(key, vals)` is run for each unique key emitted by `map()`
 - ▶ emits final output

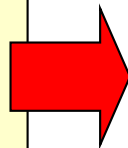
count words in docs

- ▶ Input consists of (url, contents) pairs
- ▶ map(key = url, val = contents):
 - ▶ For each word w in contents, emit (w, “1”)
- ▶ reduce(key = word, values = uniq_counts):
 - ▶ Sum all “1”s in values list
 - ▶ Emit result “(word, sum)”


Count, Illustrated

- ▶ `map(key=url, val=contents):`
 - ▶ For each word `w` in `contents`, emit `(w, "1")`
- ▶ `reduce(key=word, values=uniq_counts):`
 - ▶ Sum all "1"s in values list
 - ▶ Emit result `"(word, sum)"`

see bob throw
see spot run



see 1
bob 1
run 1
see 1
spot 1
throw 1



bob 1
run 1
see 2
spot 1
throw 1

Grep

- ▶ Input consists of (url+offset, single line)
- ▶ map(key=url+offset, val=line):
 - ▶ If contents matches regexp, emit (line, "1")
- ▶ reduce(key=line, values=uniq_counts):
 - ▶ Don't do anything; just emit line

Reverse Web-Link Graph

- ▶ **Map**
 - ▶ For each URL linking to target, ...
 - ▶ Output $\langle \text{target}, \text{source} \rangle$ pairs
- ▶ **Reduce**
 - ▶ Concatenate list of all source URLs
 - ▶ Outputs: $\langle \text{target}, \text{list}(\text{source}) \rangle$ pairs

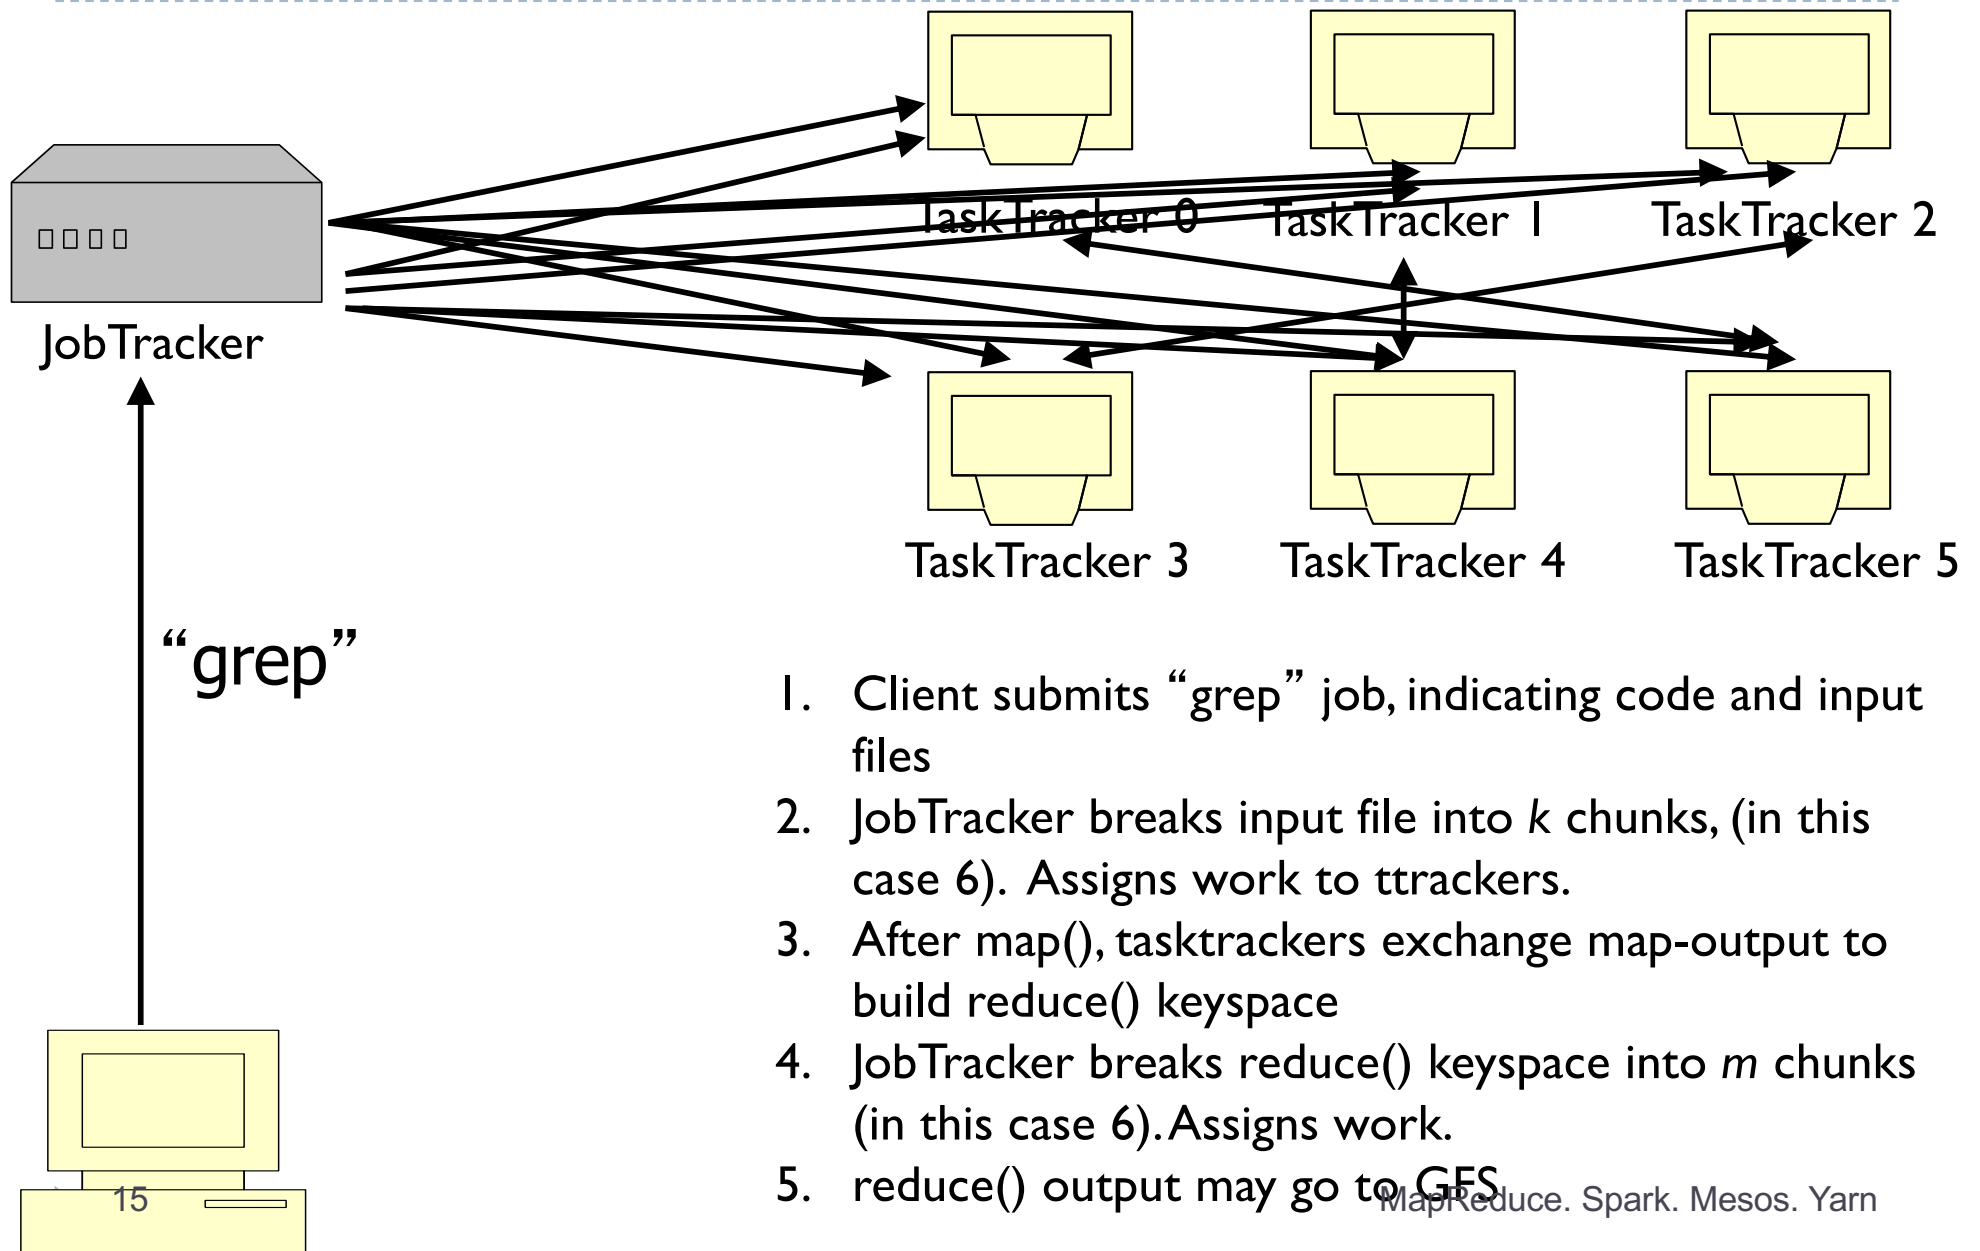
Implementation

- ▶ **Typical cluster:**
 - ▶ 100s/1000s of 2-CPU x86 machines, 2-4 GB of memory
 - ▶ Limited bisection bandwidth
 - ▶ Storage is on local IDE disks
 - ▶ GFS: distributed file system manages data
- ▶ Implementation is a C++ library linked into user programs
- ▶ Run-time system:
 - ▶ partitions the input data
 - ▶ schedules the program's execution across a set of machines
 - ▶ handles machine failures
 - ▶ manages inter-machine communication

Execution

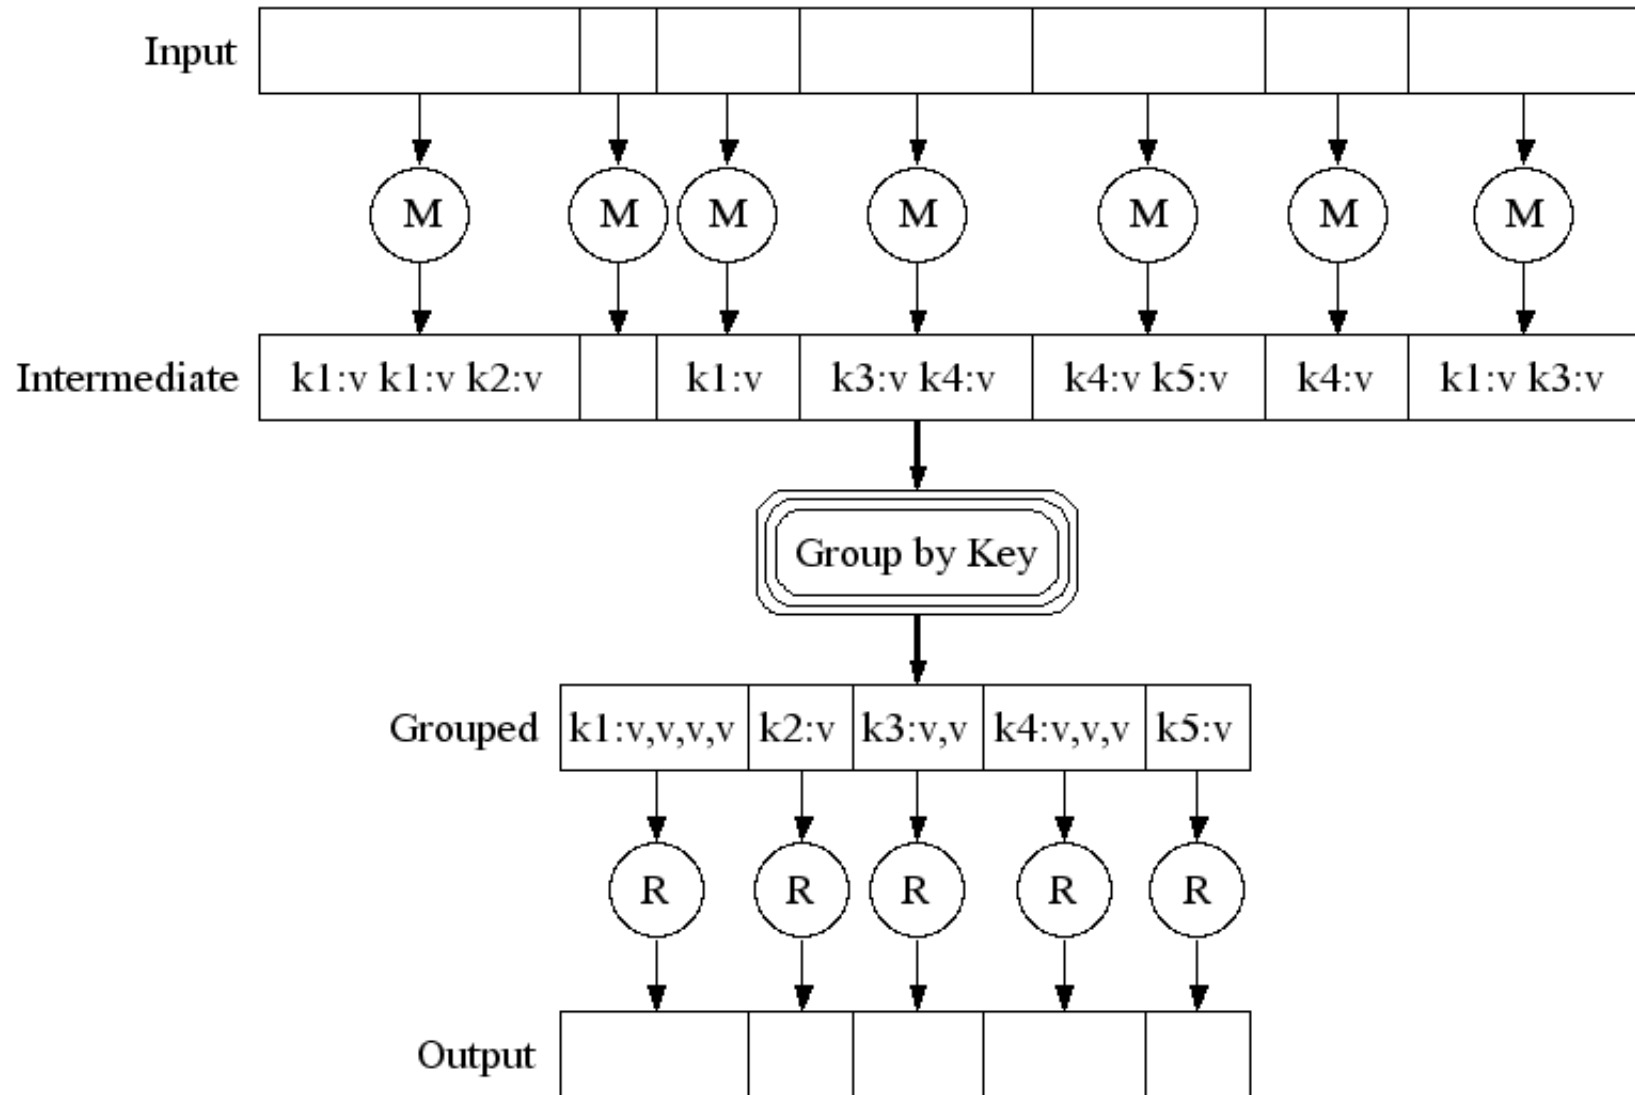
- ▶ **How is this distributed?**
 - ▶ Partition input key/value pairs into chunks, run `map()` tasks in parallel
 - ▶ After all `map()`s are complete, consolidate all emitted values for each unique emitted key
 - ▶ Now partition space of output map keys, and run `reduce()` in parallel
- ▶ **If `map()` or `reduce()` fails, reexecute!**

Job Processing

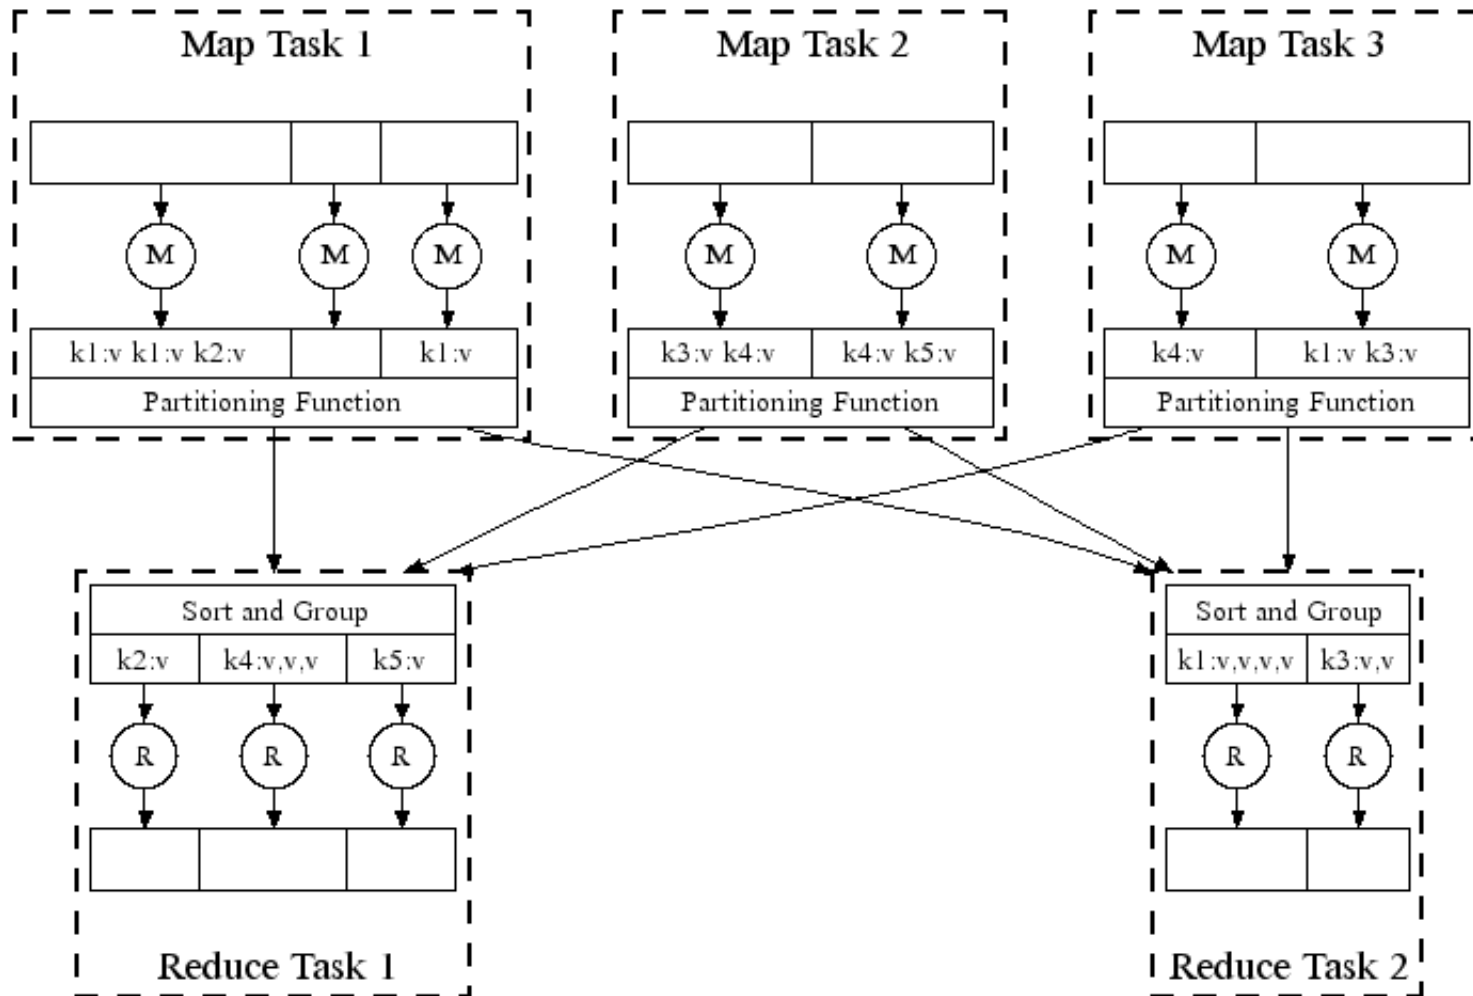


1. Client submits "grep" job, indicating code and input files
2. JobTracker breaks input file into k chunks, (in this case 6). Assigns work to ttrackers.
3. After map(), tasktrackers exchange map-output to build reduce() keyspace
4. JobTracker breaks reduce() keyspace into m chunks (in this case 6). Assigns work.
5. reduce() output may go to GFS

Execution

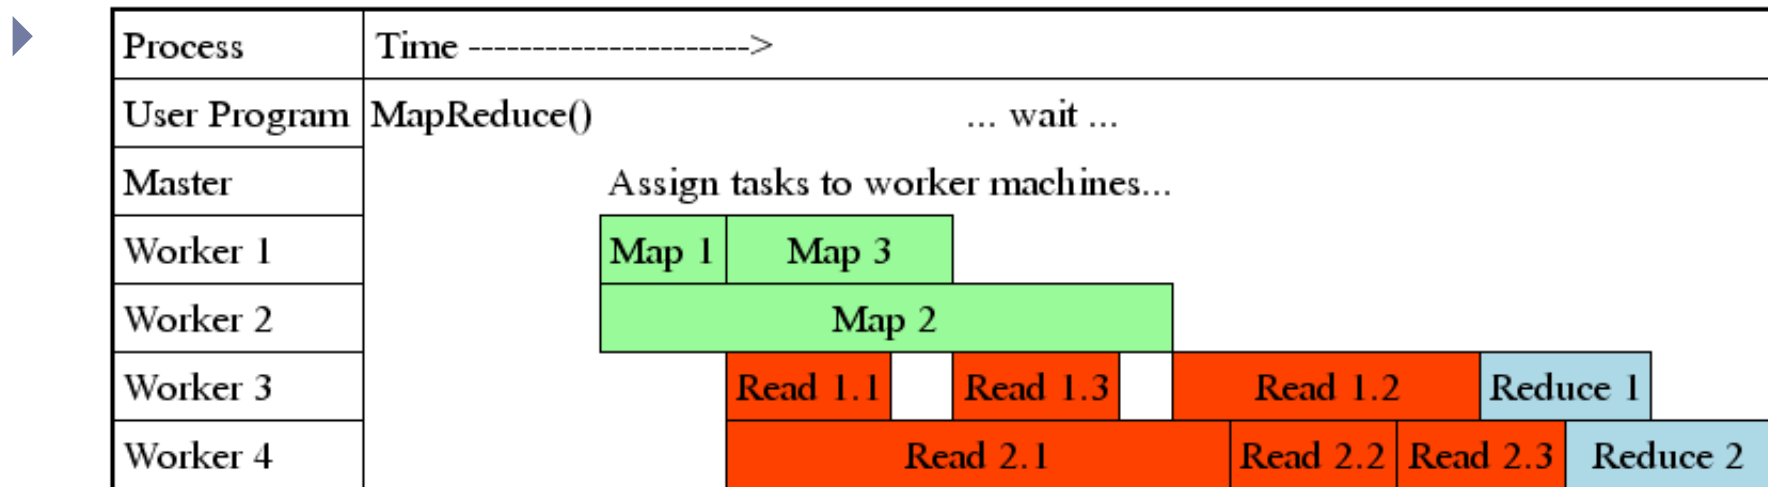


Parallel Execution



Task Granularity & Pipelining

- ▶ **Fine granularity tasks: map tasks \gg machines**
 - ▶ Minimizes time for fault recovery
 - ▶ Can pipeline shuffling with map execution
 - ▶ Better dynamic load balancing
- ▶ **Often use 200,000 map & 5000 reduce tasks**



Fault Tolerance / Workers

- ▶ **Handled via re-execution**
 - ▶ Detect failure via periodic heartbeats
 - ▶ Re-execute completed + in-progress map tasks
 - ▶ Re-execute in progress reduce tasks
 - ▶ Task completion committed through master
- ▶ **Robust: lost 1600/1800 machines once → finished ok**

Master Failure

- ▶ **Master keeps several data structures.**
 - ▶ For each map task and reduce task, it stores the state (idle, in-progress, or completed), and the identity of the worker machine (for non-idle tasks).
 - ▶ For each completed map task, the master stores the locations and sizes of the R intermediate file regions produced by the map task. The information is pushed incrementally to workers that have in-progress reduce tasks.
- ▶ **There is no fault-tolerance for the master**
 - ▶ Can be done with periodic checkpoints and starting a new copy from the last checkpointed state
 - ▶ Current implementation aborts the MapReduce computation if the master fails.

Refinement: Redundant Execution

Slow workers significantly delay completion time

- ▶ Other jobs consuming resources on machine
- ▶ Bad disks w/ soft errors transfer data slowly
- ▶ Weird things: processor caches disabled (!!)

Solution: Near end of phase, spawn backup tasks

- ▶ Whichever one finishes first "wins"

Dramatically shortens job completion time

Refinement: Locality Optimization

- ▶ **Master scheduling policy:**
 - ▶ Asks GFS for locations of replicas of input file blocks
 - ▶ Map tasks typically split into 64MB (GFS block size)
 - ▶ Map tasks scheduled so GFS input block replica are on same machine or same rack

- ▶ **Effect**
 - ▶ Thousands of machines read input at local disk speed
 - ▶ Without this, rack switches limit read rate

Refinement: Skipping Bad Records

- ▶ **Map/Reduce functions sometimes fail for particular inputs**
 - ▶ Best solution is to debug & fix
 - ▶ Not always possible ~ third-party source libraries
 - ▶ On segmentation fault:
 - ▶ Send UDP packet to master from signal handler
 - ▶ Include sequence number of record being processed
 - ▶ If master sees two failures for same record:
 - ▶ Next worker is told to skip the record

Other Refinements

- ▶ **Sorting guarantees**
 - ▶ within each reduce partition
- ▶ **Compression of intermediate data**
- ▶ **Combiner**
 - ▶ useful for saving network bandwidth
- ▶ **Local execution for debugging/testing**
- ▶ **User-defined counters**

Performance

Tests run on cluster of 1800 machines:

- ▶ 4 GB of memory
- ▶ Dual-processor 2 GHz Xeons with Hyperthreading
- ▶ Dual 160 GB IDE disks
- ▶ Gigabit Ethernet per machine
- ▶ Bisection bandwidth approximately 100 Gbps

Two benchmarks:

MR_GrepScan 1010 100-byte records to extract records matching a rare pattern (92K matching records)

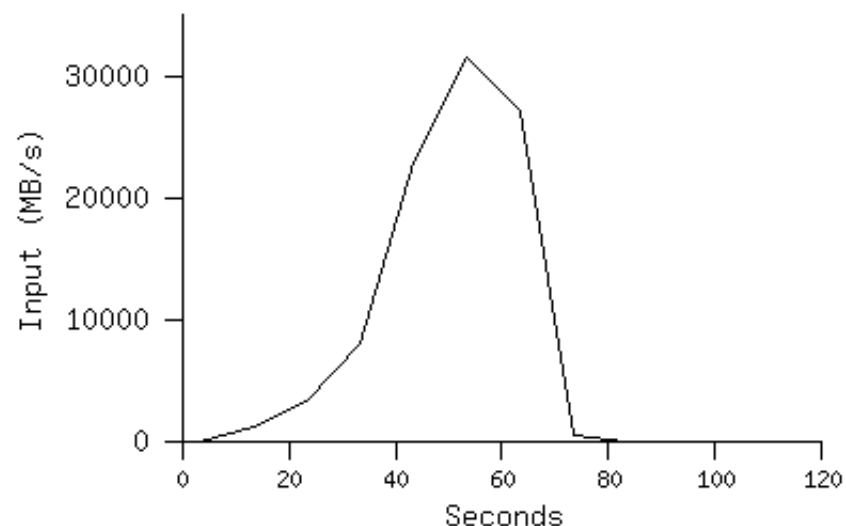
▶ **MR_SortSort** 1010 100-byte records (modeled after TeraSort benchmark)

MR_Grep

Locality optimization helps:

- ▶ 1800 machines read 1 TB at peak ~31 GB/s
- ▶ W/out this, rack switches would limit to 10 GB/s

Startup overhead is significant for short jobs

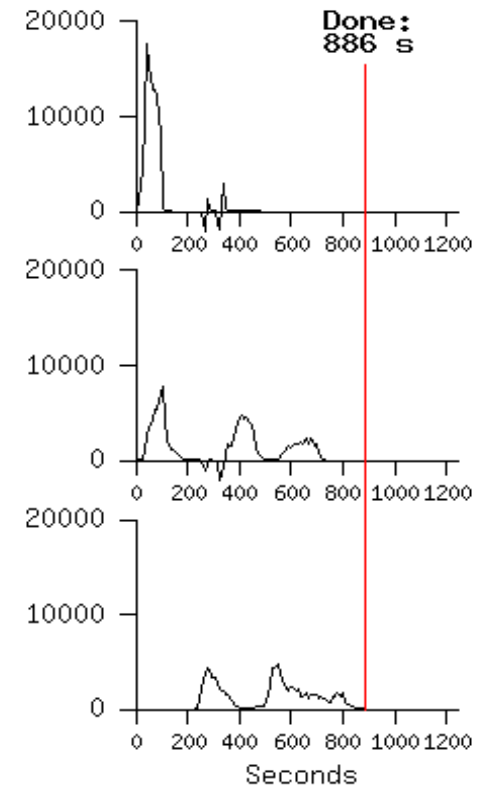
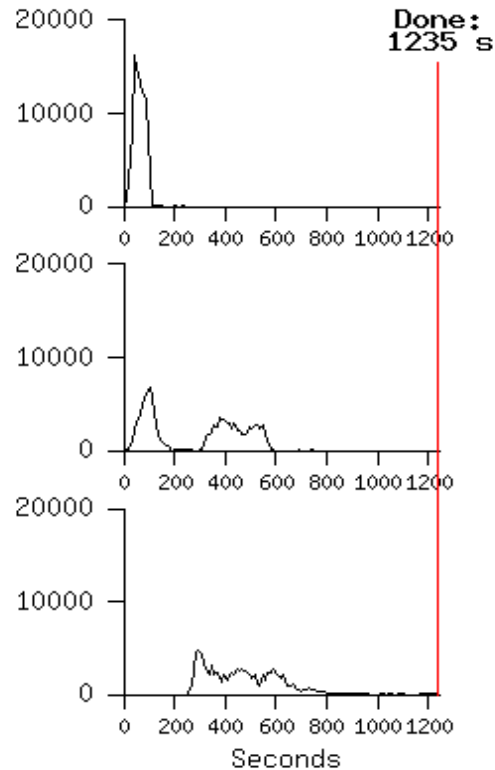
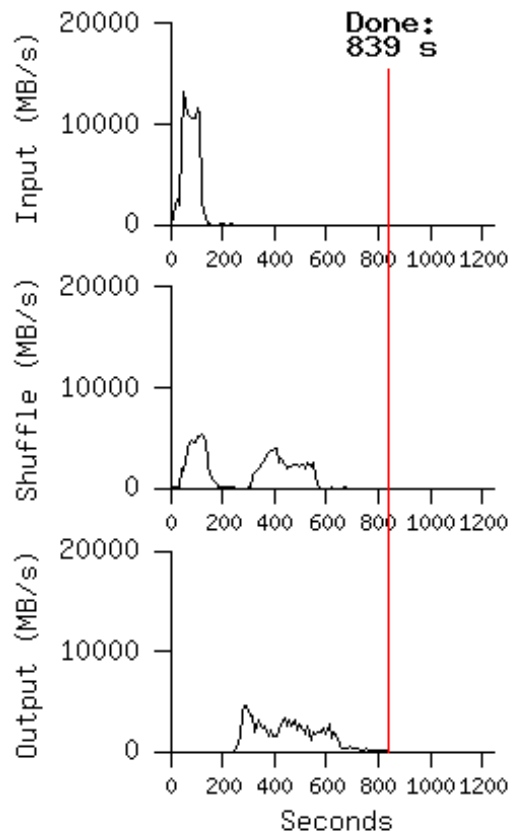


MR_Sort

Normal

No backup tasks

200 processes killed




- Backup tasks reduce job completion time a lot!
- System deals well with failures



2: Hadoop

Apache Hadoop

hadoop

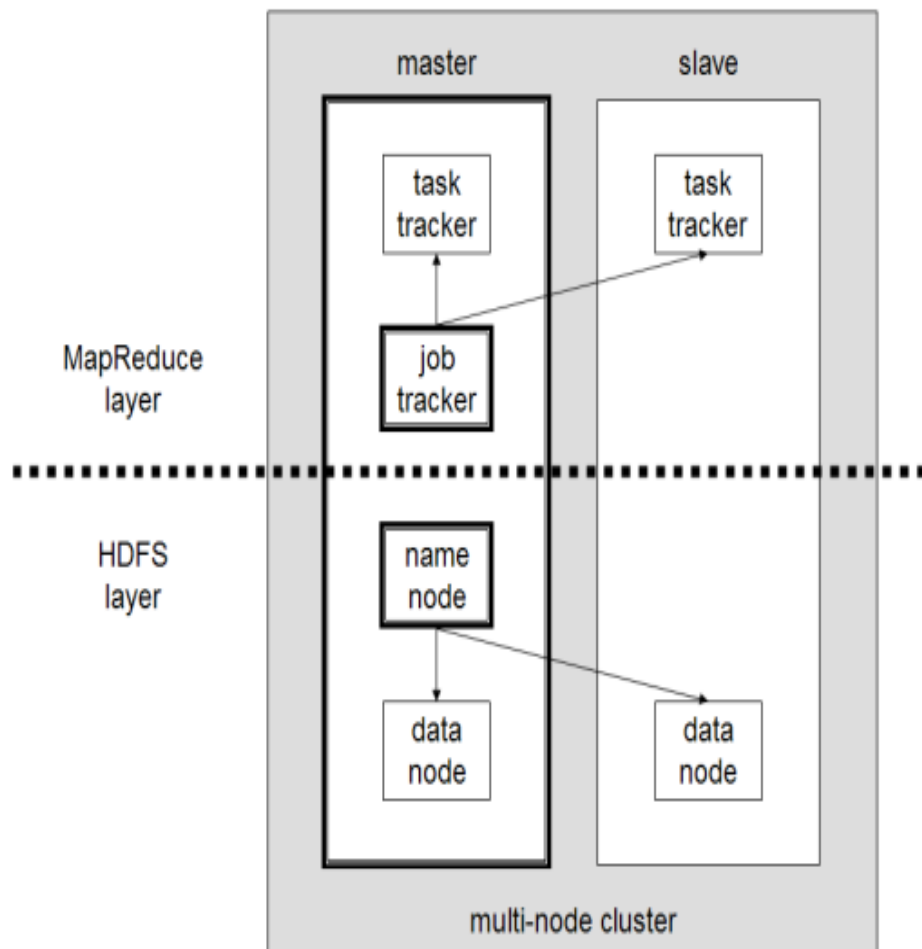


- ▶ Apache Hadoop's MapReduce and HDFS components originally derived from
 - ▶ Google File System (GFS)¹ – 2003
 - ▶ Google's MapReduce² - 2004
- ▶ Data is broken in splits that are processed in different machines.
- ▶ Industry wide standard for processing Big Data.

Overview of Hadoop

- ▶ **Basic components of Hadoop are:**
 - ▶ **Map Reduce Layer**
 - ▶ **Job tracker** (master) -which coordinates the execution of jobs;
 - ▶ **Task trackers** (slaves)- which control the execution of map and reduce tasks in the machines that do the processing;
 - ▶ **HDFS Layer-** which stores files.
 - ▶ **Name Node** (master)- manages the file system, keeps metadata for all the files and directories in the tree
 - ▶ **Data Nodes** (slaves)- work horses of the file system. Store and retrieve blocks when they are told to (by clients or name node) and report back to name node periodically

Overview of Hadoop contd.



Job Tracker - coordinates the execution of jobs

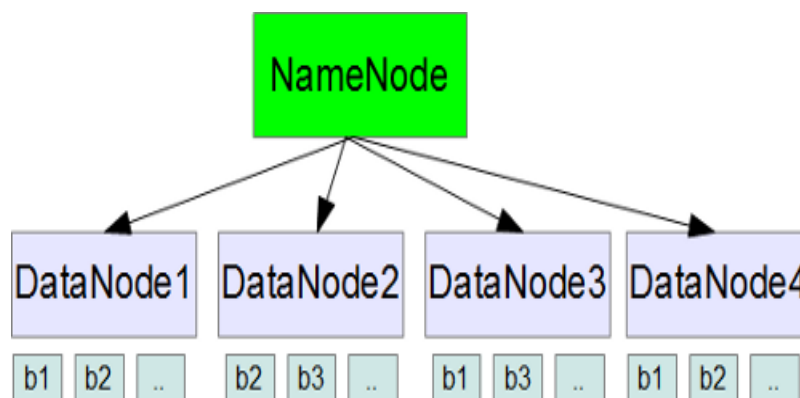
Task Tracker – control the execution of map and reduce tasks in slave machines

Name Node – Manages the file system, keeps metadata

Data Node – Follow the instructions from name node
- stores, retrieves data

Fault Tolerance in HDFS layer

- ▶ Hardware failure is the norm rather than the exception
- ▶ **Detection of faults and quick, automatic recovery from them** is a core architectural goal of HDFS.
- ▶ Master Slave Architecture with NameNode (master) and DataNode (slave)
- ▶ Common types of failures
 - ▶ NameNode failures
 - ▶ DataNode failures



Handling Data Node Failure

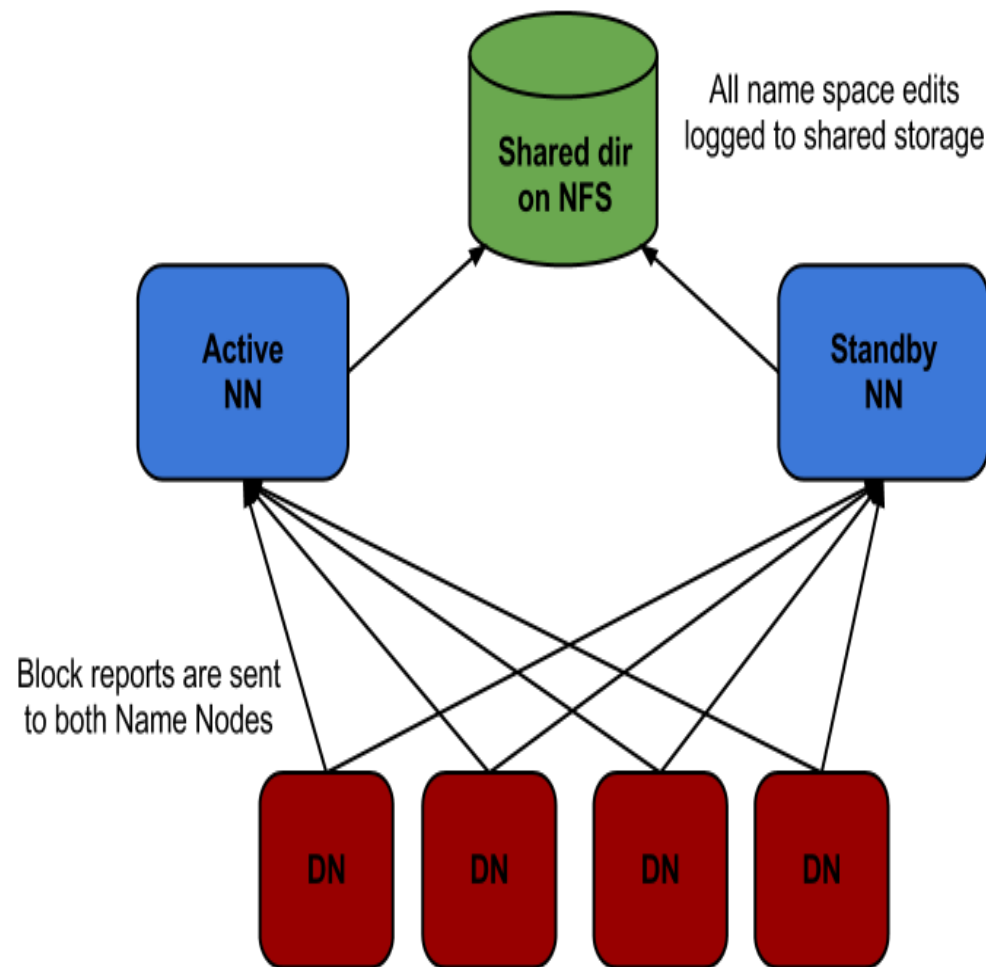
- ▶ Each DataNode sends a Heartbeat message to the NameNode periodically
- ▶ If the namenode does not receive a heartbeat from a particular data node for 10 minutes, then it considers that data node to be dead/out of service.
- ▶ Name Node initiates replication of blocks which were hosted on that data node to be hosted on some other data node.

Handling Name Node Failure

- ▶ Single Name Node per cluster.
- ▶ Prior to Hadoop 2.0.0, the NameNode was a single point of failure (SPOF) in an HDFS cluster.
- ▶ If NameNode becomes unavailable, the cluster as a whole would be unavailable
 - ▶ NameNode has to be restarted
 - ▶ Brought up on a separate machine.

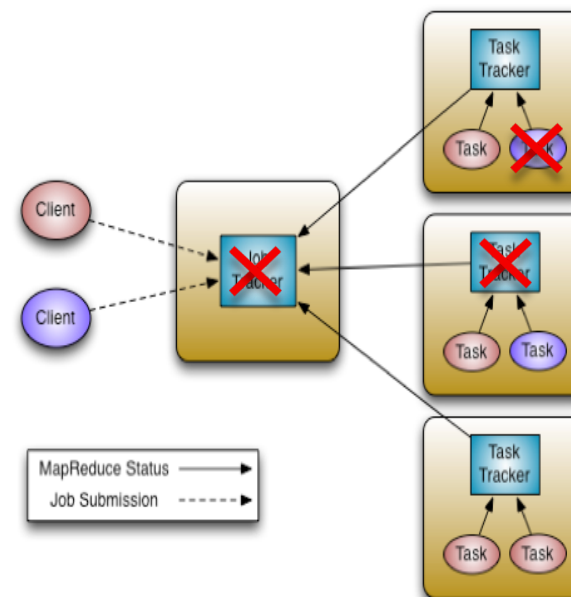
HDFS High Availability

- ▶ Provides an option of running two redundant NameNodes in the same cluster
- ▶ Active/Passive configuration with a hot standby.
- ▶ Fast failover to a new NameNode in the case that a machine crashes
- ▶ Graceful administrator-initiated failover for the purpose of planned maintenance.



Classic MapReduce (v1)

- ▶ **Job Tracker**
 - ▶ Manage Cluster Resources and Job Scheduling
- ▶ **Task Tracker**
 - ▶ Per-node agent
 - ▶ Manage Tasks
- ▶ **Jobs can fail**
 - ▶ While running the task (Task Failure)
 - ▶ Task Tracker failure
 - ▶ Job Tracker failure



Handling Task Failure

- ▶ **User code bug in map/reduce**
 - ▶ Throws a `RuntimeException`
 - ▶ Child JVM reports a failure back to the parent task tracker before it exits.
- ▶ **Sudden exit of the child JVM**
 - ▶ Bug that causes the JVM to exit for the conditions exposed by map/reduce code.
- ▶ Task tracker marks the task attempt as failed, makes room available to another task.

Task Tracker Failure

- ▶ Task tracker will stop sending the heartbeat to the Job Tracker
- ▶ Job Tracker notices this failure
 - ▶ Hasn't received a heart beat from 10 mins
 - ▶ Can be configured via `mapred.tasktracker.expiry.interval` property
- ▶ Job Tracker removes this task from the task pool
- ▶ Rerun the Job even if map task has ran completely
 - ▶ Intermediate output resides in the failed task trackers local file system which is not accessible by the reduce tasks.

Job Tracker Failure

- ▶ This is serious than the other two modes of failure.
 - ▶ Single point of failure.
 - ▶ In this case all jobs will fail.
- ▶ After restarting Job Tracker all the jobs running at the time of the failure needs to be resubmitted.



3. Spark

Slides by Matei Zaharia, UC Berkeley

Motivation

- ▶ **Map reduce based tasks are slow**
 - ▶ Sharing of data across jobs is stable storage
 - ▶ Replication of data and disk I/O
- ▶ **Support iterative algorithms**
- ▶ **Support interactive data mining tools – search**



Existing literature on large distributed algorithms on clusters

- ▶ **General** : Language-integrated “distributed dataset” API, but cannot share datasets efficiently across queries
 - ▶ Map Reduce
 - ▶ Map
 - ▶ Shuffle
 - ▶ Reduce
 - ▶ DyradLinq
 - ▶ Ciel
- ▶ **Specific** : Specialized models; can't run arbitrary / ad-hoc queries
 - ▶ Pregel – Google's graph based
 - ▶ Haloop – iterative Hadoop
- ▶

▶ Caching systems

- ▶ Nectar – Automatic expression caching, but over distributed FS
- ▶ Ciel – not explicit control over cached data
- ▶ PacMan - Memory cache for HDFS, but writes still go to network/disk

▶ Lineage

- ▶ To track dependency of task information across a DAG of tasks



What is Spark?

- ▶ Fast, expressive cluster computing system compatible with Apache Hadoop
 - ▶ Works with any Hadoop-supported storage system (HDFS, S3, Avro, ...)
- ▶ Improves **efficiency** through:
 - ▶ In-memory computing primitives
 - ▶ General computation graphs
- ▶ Improves **usability** through:
 - ▶ Rich APIs in Java, Scala, Python
 - ▶ Interactive shell

Key Idea

- ▶ **Work with distributed collections as you would with local ones**
- ▶ **Concept: resilient distributed datasets (RDDs)**
 - ▶ Immutable collections of objects spread across a cluster
 - ▶ Built through parallel transformations (map, filter, etc)
 - ▶ Automatically rebuilt on failure
 - ▶ Controllable persistence (e.g. caching in RAM)

Resilient Distributed Datasets (RDDs)

- ▶ **Restricted form of distributed shared memory**
 - ▶ Read only/ Immutable , partitioned collections of records
 - ▶ Deterministic
 - ▶ From coarse grained operations (map, filter, join, etc.)
 - ▶ From stable storage or other RDDs
 - ▶ User controlled persistence
 - ▶ User controlled partitioning



Spark programming interface

- ▶ **Lazy operations**
 - ▶ Transformations not done until action
- ▶ **Operations on RDDs**
 - ▶ Transformations - build new RDDs
 - ▶ Actions - compute and output results
- ▶ **Partitioning – layout across nodes**
- ▶ **Persistence – storage in RAM / Disc**



Representing RDDs

Operation	Meaning
<code>partitions()</code>	Return a list of Partition objects
<code>preferredLocations(<i>p</i>)</code>	List nodes where partition <i>p</i> can be accessed faster due to data locality
<code>dependencies()</code>	Return a list of dependencies
<code>iterator(<i>p</i>, <i>parentIters</i>)</code>	Compute the elements of partition <i>p</i> given iterators for its parent partitions
<code>partitioner()</code>	Return metadata specifying whether the RDD is hash/range partitioned

Table 3: Interface used to represent RDDs in Spark.

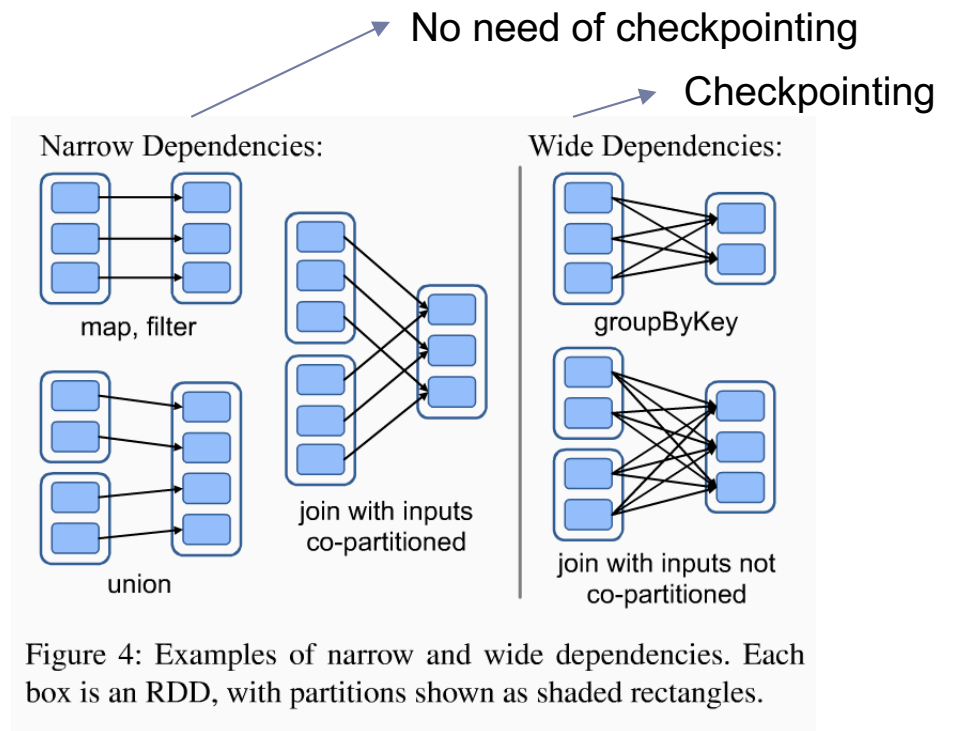


Figure 4: Examples of narrow and wide dependencies. Each box is an RDD, with partitions shown as shaded rectangles.

RDD on Spark

Transformations	<ul style="list-style-type: none"><i>map</i>($f : T \Rightarrow U$) : $RDD[T] \Rightarrow RDD[U]$<i>filter</i>($f : T \Rightarrow Bool$) : $RDD[T] \Rightarrow RDD[T]$<i>flatMap</i>($f : T \Rightarrow Seq[U]$) : $RDD[T] \Rightarrow RDD[U]$<i>sample</i>(<i>fraction</i> : Float) : $RDD[T] \Rightarrow RDD[T]$ (Deterministic sampling)<i>groupByKey</i>() : $RDD[(K, V)] \Rightarrow RDD[(K, Seq[V])]$<i>reduceByKey</i>($f : (V, V) \Rightarrow V$) : $RDD[(K, V)] \Rightarrow RDD[(K, V)]$<i>union</i>() : $(RDD[T], RDD[T]) \Rightarrow RDD[T]$<i>join</i>() : $(RDD[(K, V)], RDD[(K, W)]) \Rightarrow RDD[(K, (V, W))]$<i>cogroup</i>() : $(RDD[(K, V)], RDD[(K, W)]) \Rightarrow RDD[(K, (Seq[V], Seq[W]))]$<i>crossProduct</i>() : $(RDD[T], RDD[U]) \Rightarrow RDD[(T, U)]$<i>mapValues</i>($f : V \Rightarrow W$) : $RDD[(K, V)] \Rightarrow RDD[(K, W)]$ (Preserves partitioning)<i>sort</i>($c : Comparator[K]$) : $RDD[(K, V)] \Rightarrow RDD[(K, V)]$<i>partitionBy</i>($p : Partitioner[K]$) : $RDD[(K, V)] \Rightarrow RDD[(K, V)]$
Actions	<ul style="list-style-type: none"><i>count</i>() : $RDD[T] \Rightarrow Long$<i>collect</i>() : $RDD[T] \Rightarrow Seq[T]$<i>reduce</i>($f : (T, T) \Rightarrow T$) : $RDD[T] \Rightarrow T$<i>lookup</i>($k : K$) : $RDD[(K, V)] \Rightarrow Seq[V]$ (On hash/range partitioned RDDs)<i>save</i>(<i>path</i> : String) : Outputs RDD to a storage system, <i>e.g.</i>, HDFS

Table 2: Transformations and actions available on RDDs in Spark. $Seq[T]$ denotes a sequence of elements of type T.



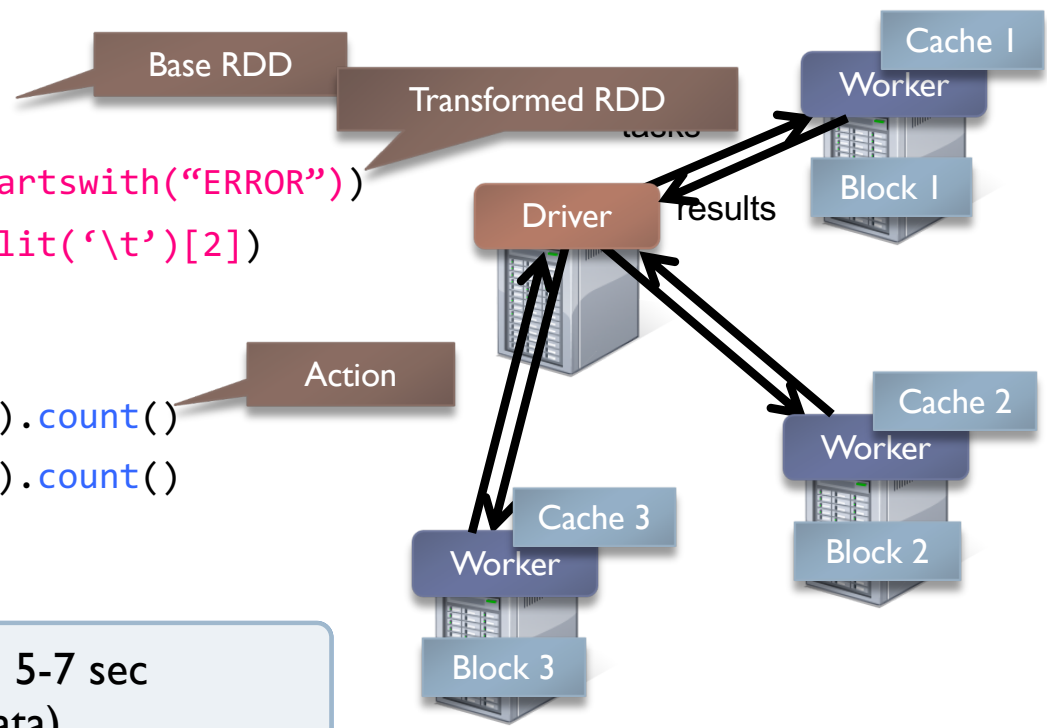
Example: Mining Console Logs

- ▶ Load error messages from a log into memory, then interactively search for patterns

```
lines = spark.textFile("hdfs://...")
errors = lines.filter(lambda s: s.startswith("ERROR"))
messages = errors.map(lambda s: s.split('\t')[2])
messages.cache()

messages.filter(lambda s: "foo" in s).count()
messages.filter(lambda s: "bar" in s).count()
. . .
```

Result: scaled to 1 TB data in 5-7 sec
(vs 170 sec for on-disk data)



Example : Console Log mining

```
lines = spark.textFile("hdfs://...")
errors = lines.filter(_.startsWith("ERROR"))
errors.persist()

errors.count()

// Count errors mentioning MySQL:
errors.filter(_.contains("MySQL")).count()

// Return the time fields of errors mentioning
// HDFS as an array (assuming time is field
// number 3 in a tab-separated format):
errors.filter(_.contains("HDFS"))
    .map(_.split('\t')(3))
    .collect()
```

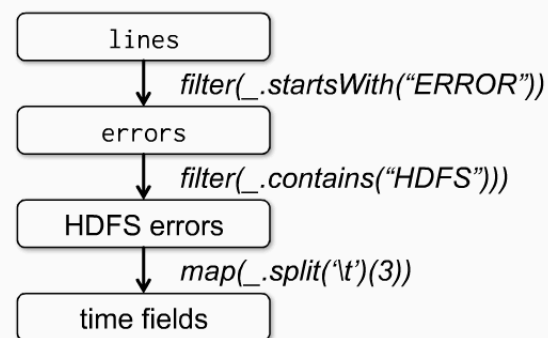


Figure 1: Lineage graph for the third query in our example. Boxes represent RDDs and arrows represent transformations.

Example : Logistic regression

- ▶ Classification problem that searches for hyper plane w

```
val points = spark.textFile(...)
                    .map(parsePoint).persist()
var w = // random initial vector
for (i <- 1 to ITERATIONS) {
  val gradient = points.map{ p =>
    p.x * (1/(1+exp(-p.y*(w dot p.x)))-1)*p.y
  }.reduce((a,b) => a+b)
  w -= gradient
}
```

Transforms text to point object

Repetitive map and reduce
to compute gradient



Example : PageRank

- ▶ Start each page with rank $1/N$.
- ▶ On each iteration update the page rank
 - ▶ $= \sum_{i \in \text{neighbors}} \text{rank}_i / |\text{neighbors}|$

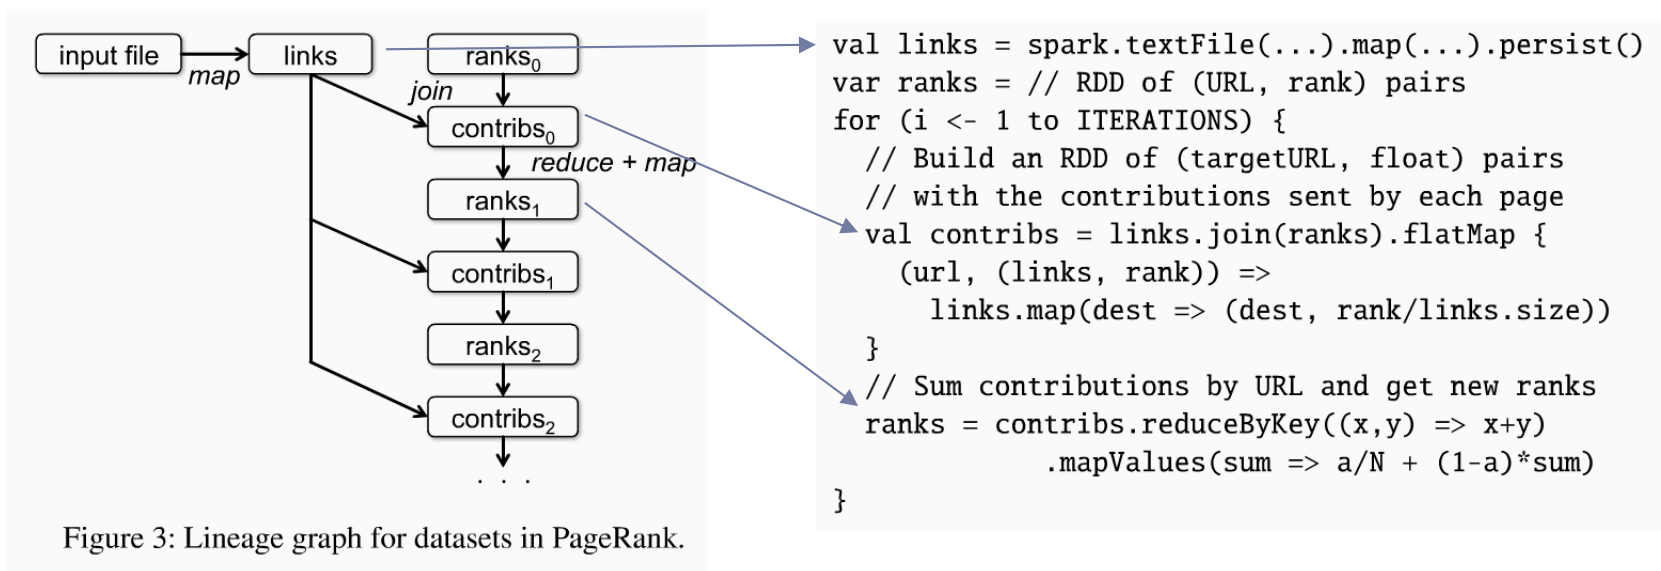
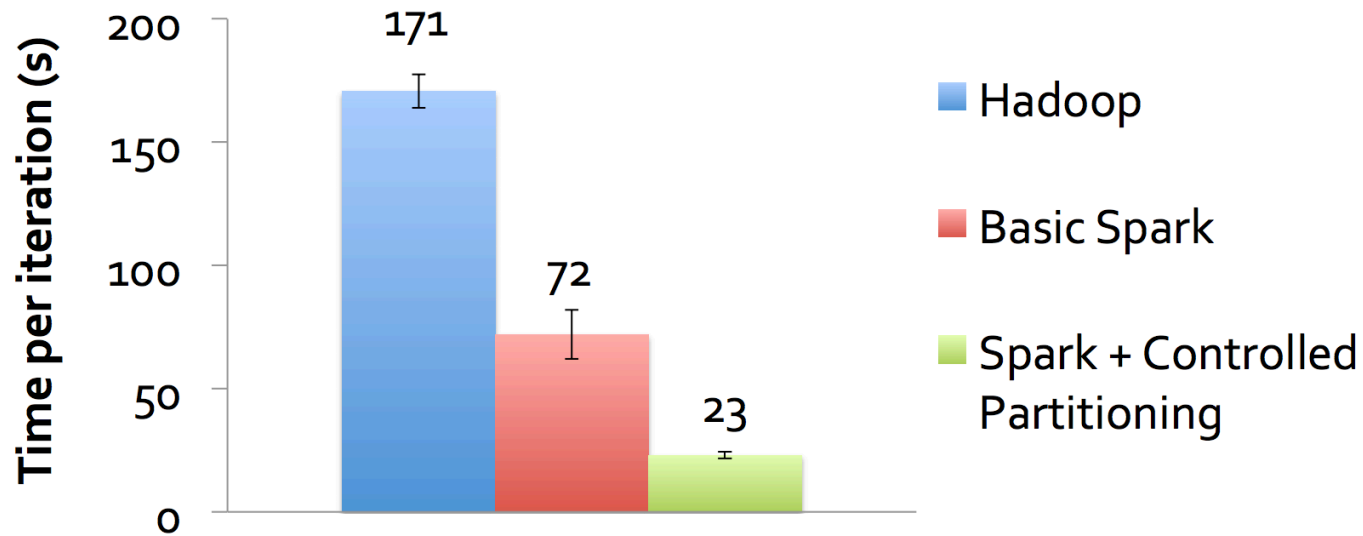


Figure 3: Lineage graph for datasets in PageRank.

PageRank performance



RDDs versus DSMs

Lookup by key



Aspect	RDDs	Distr. Shared Mem.
Reads	Coarse- or fine-grained	Fine-grained
Writes	Coarse-grained	Fine-grained
Consistency	Trivial (immutable)	Up to app / runtime
Fault recovery	Fine-grained and low-overhead using lineage	Requires checkpoints and program rollback
Straggler mitigation	Possible using backup tasks	Difficult
Work placement	Automatic based on data locality	Up to app (runtimes aim for transparency)
Behavior if not enough RAM	Similar to existing data flow systems	Poor performance (swapping?)

Table 1: Comparison of RDDs with distributed shared memory.

RDDs unsuitable for applications that make fine-grained updates to shared state, -storage system for a web application -an incremental web crawler



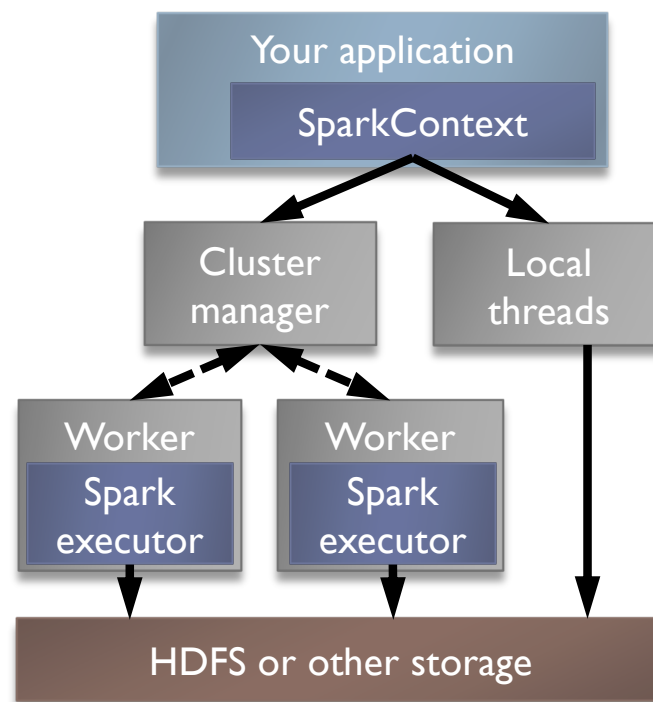
Implementation in Spark

- ▶ **Job scheduler**
 - ▶ Data locality captured using delay scheduling
- ▶ **Interpreter integration**
 - ▶ Class shipping
 - ▶ Modified code generation
- ▶ **Memory management**
 - ▶ In memory and swap memory
 - ▶ LRU
- ▶ **Support for checkpointing**
 - ▶ Good for long lineage graphs



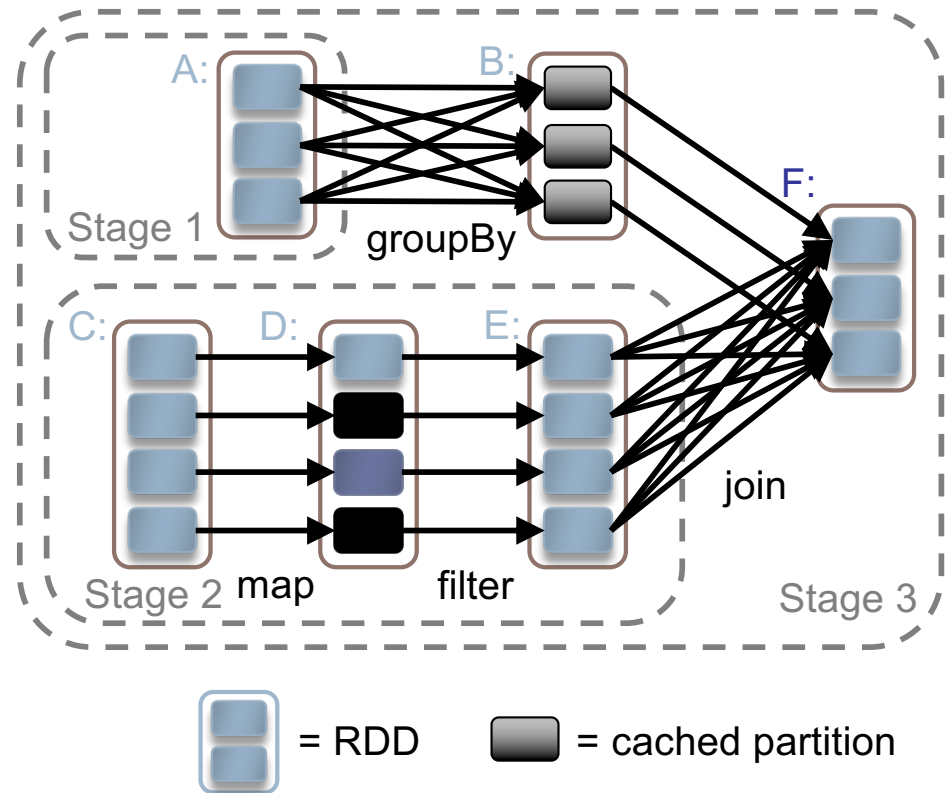
Software Components

- ▶ Spark runs as a library in your program (one instance per app)
- ▶ Runs tasks locally or on a cluster
 - ▶ Standalone deploy cluster, Mesos or YARN
- ▶ Accesses storage via Hadoop InputFormat API
 - ▶ Can use HBase, HDFS, S3, ...



Task Scheduler

- ▶ Supports general task graphs
- ▶ Pipelines functions where possible
- ▶ Cache-aware data reuse & locality
- ▶ Partitioning-aware to avoid shuffles



Hadoop Compatibility

- ▶ Spark can read/write to any storage system / format that has a plugin for Hadoop!
 - ▶ Examples: HDFS, S3, HBase, Cassandra, Avro, SequenceFile
 - ▶ Reuses Hadoop's InputFormat and OutputFormat APIs
- ▶ APIs like `sparkContext.textFile` support filesystems, while `sparkContext.hadoopRDD` allows passing any Hadoop JobConf to configure an input source

Evaluation

- ▶ Runs on Mesos to share clusters with Hadoop
- ▶ Can read from any Hadoop input source (HDFS or HBase)
- ▶ RDD implemented in Spark
 - ▶ Ability to be used over any other cluster systems as well



Iterative ML applications

scalability

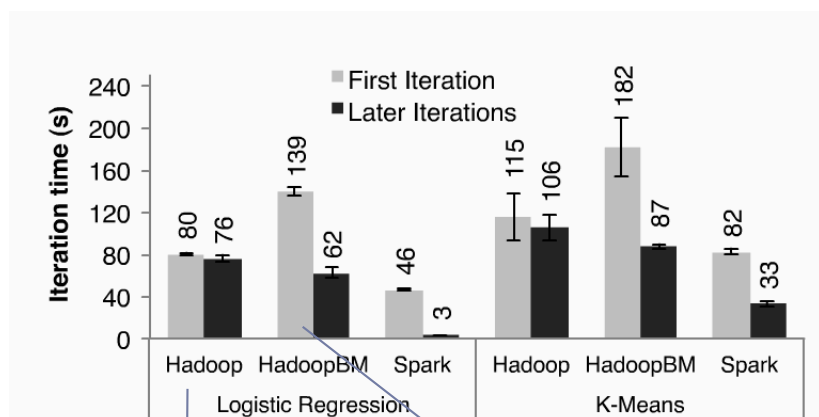
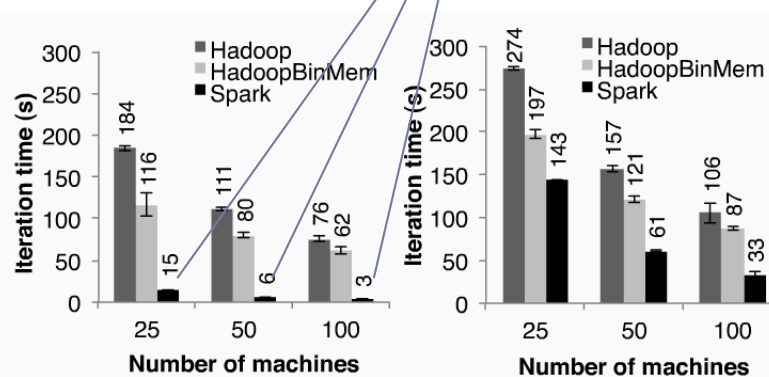


Figure 7: Duration of the first and later iterations in Hadoop, HadoopBinMem and Spark for logistic regression and k-means using 100 GB of data on a 100-node cluster.

No improvement in successive iterations
Slow due to heartbeat signals

Initially slow due to conversion of text to binary in-Mem and java objects



(a) Logistic Regression

(b) K-Means

Figure 8: Running times for iterations after the first in Hadoop, HadoopBinMem, and Spark. The jobs all processed 100 GB.



Understanding Speedup

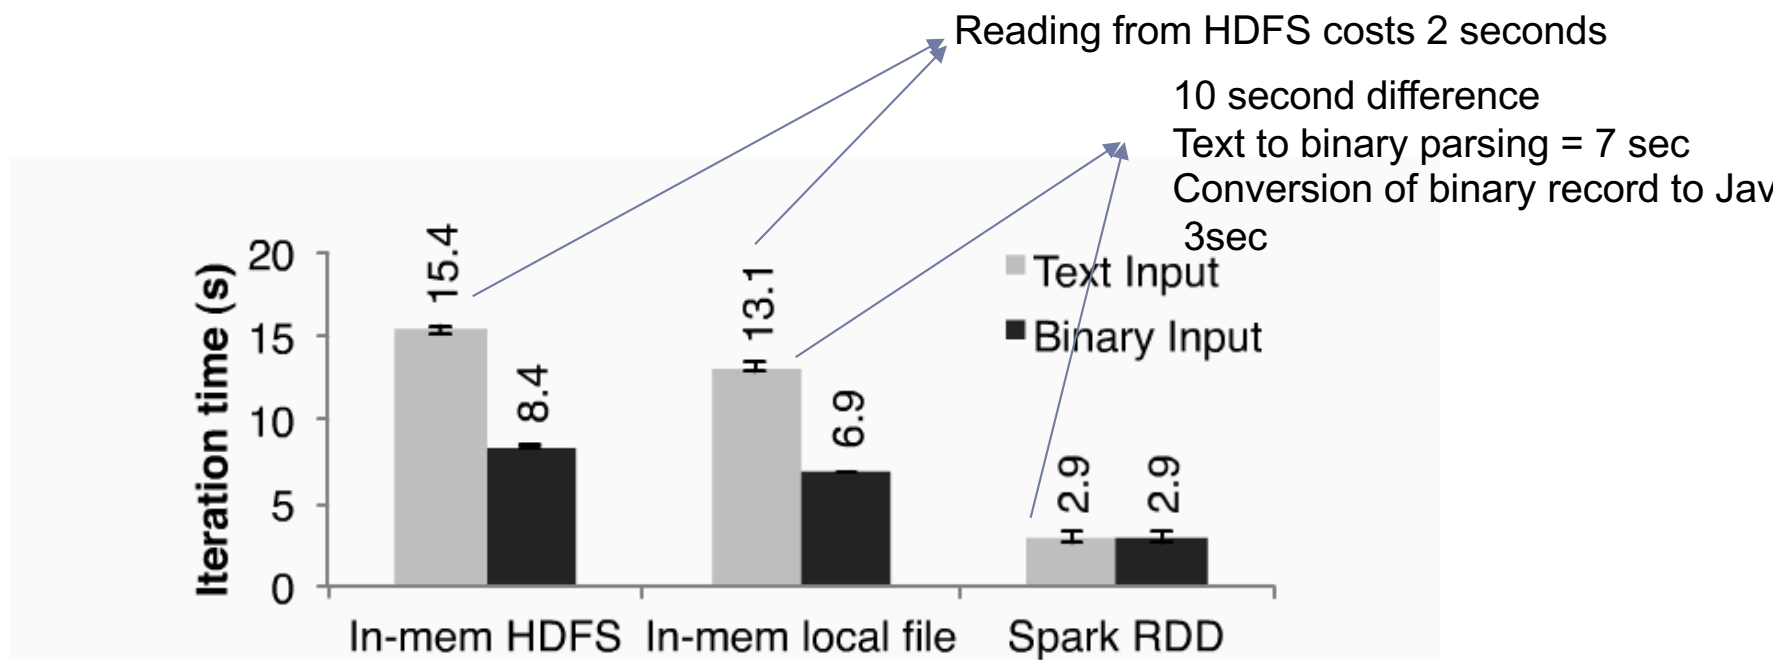
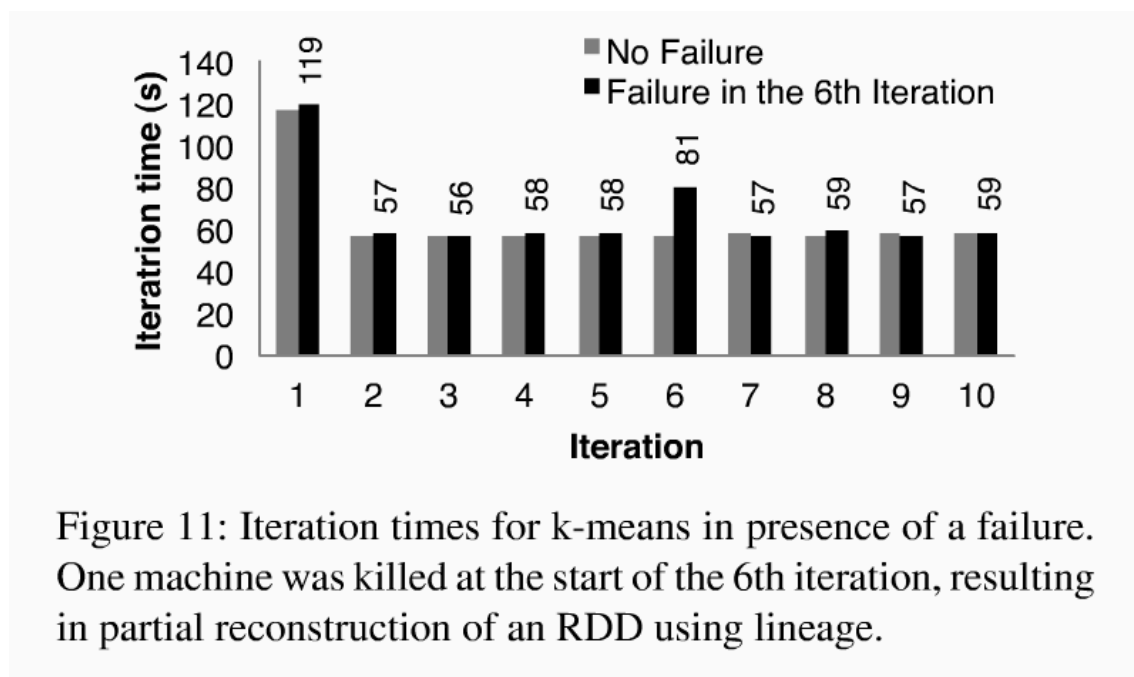


Figure 9: Iteration times for logistic regression using 256 MB data on a single machine for different sources of input.

Failure in RDD

RDDs track the graph of transformations that built them (their lineage) to rebuild lost data



Insufficient memory

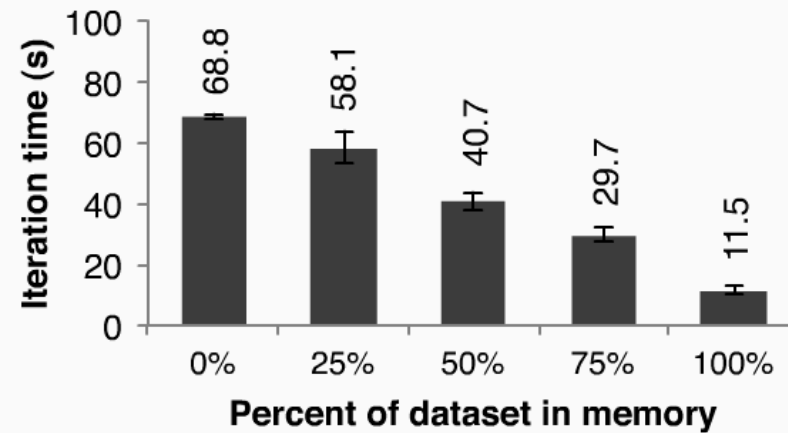
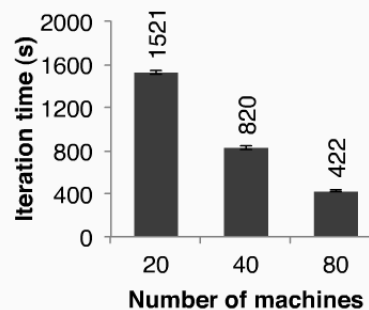


Figure 12: Performance of logistic regression using 100 GB data on 25 machines with varying amounts of data in memory.

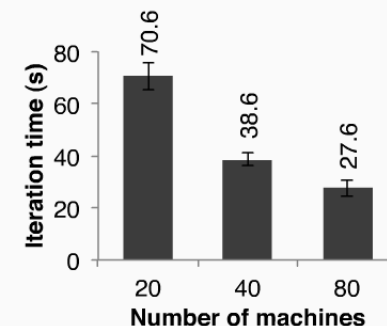


User applications using Spark

- ▶ In memory analytics at Conviva : 40x speedup
- ▶ Traffic modeling (Traffic prediction via EM - Mobile Millennium)
- ▶ Twitter spam classification(Monarch)
- ▶ DNA sequence analysis (SNAP)



(a) Traffic modeling



(b) Spam classification

Figure 13: Per-iteration running time of two user applications implemented with Spark. Error bars show standard deviations.



3: Mesos

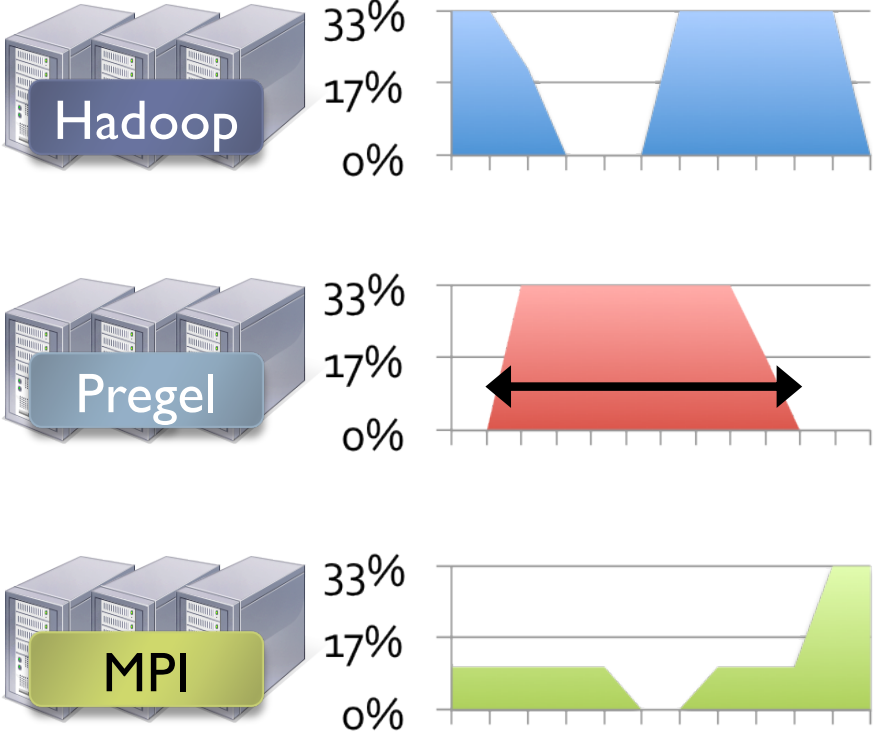
Slides by Matei Zaharia

Problem

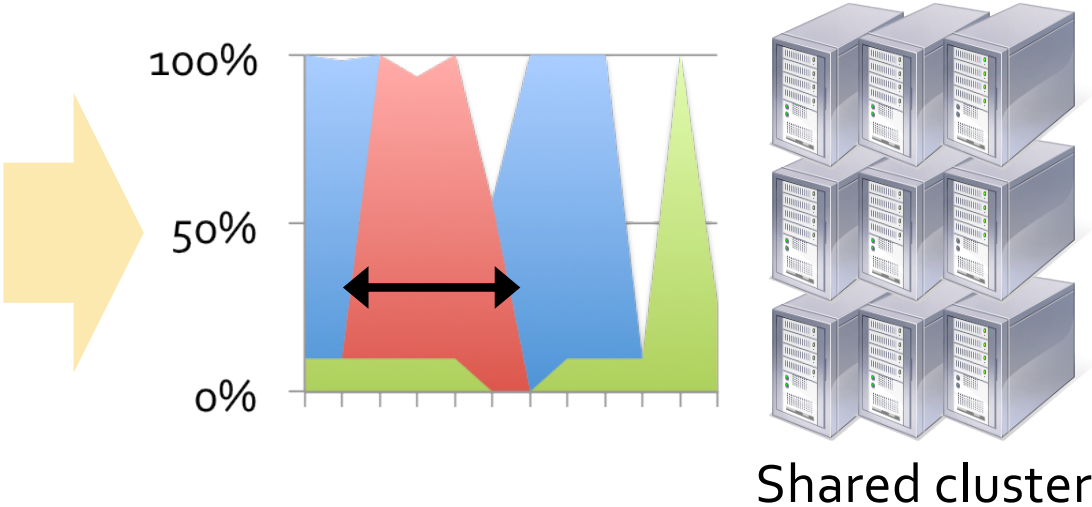
- ▶ Rapid innovation in cluster computing frameworks
- ▶ No single framework optimal for all applications
- ▶ Want to run multiple frameworks in a single cluster
 - ▶ ...to maximize utilization
 - ▶ ...to share data between frameworks

Static vs dynamic sharing

Static partitioning

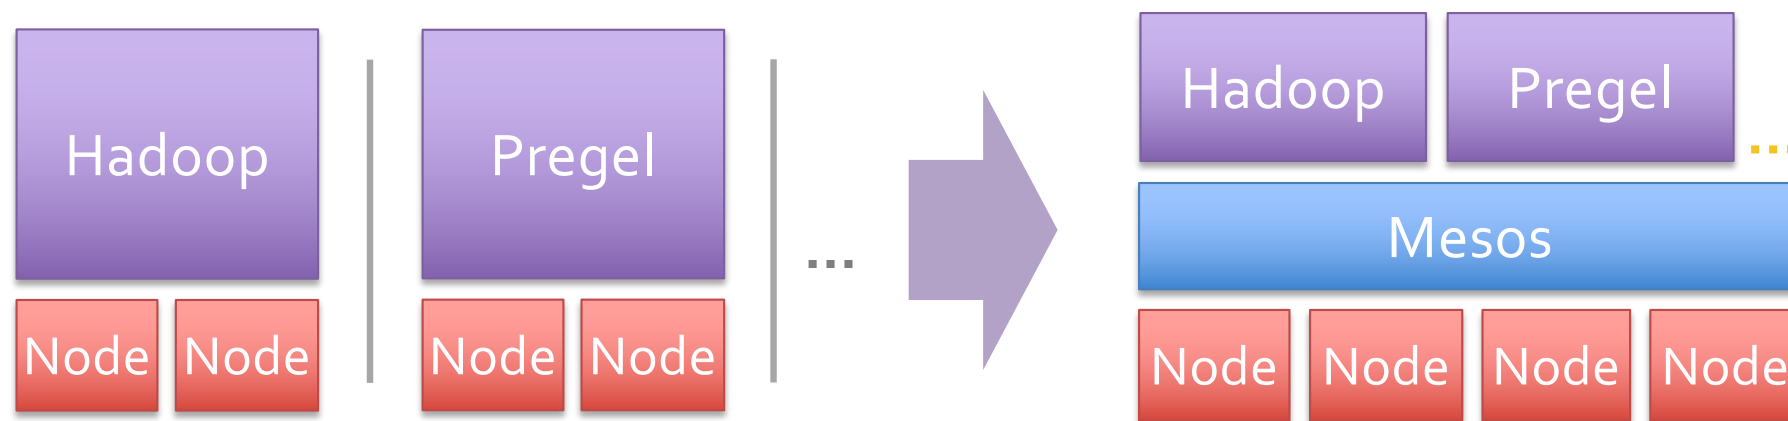


Mesos: dynamic sharing



Solution

- ▶ Mesos is a common resource sharing layer over which diverse frameworks can run



Other Benefits of Mesos

- ▶ **Run multiple instances of the same framework**
 - ▶ Isolate production and experimental jobs
 - ▶ Run multiple versions of the framework concurrently
- ▶ **Build specialized frameworks targeting particular problem domains**
 - ▶ Better performance than general-purpose abstractions

Mesos Goals

- ▶ High utilization of resources
- ▶ Support diverse frameworks (current & future)
- ▶ Scalability to 10,000's of nodes
- ▶ Reliability in face of failures

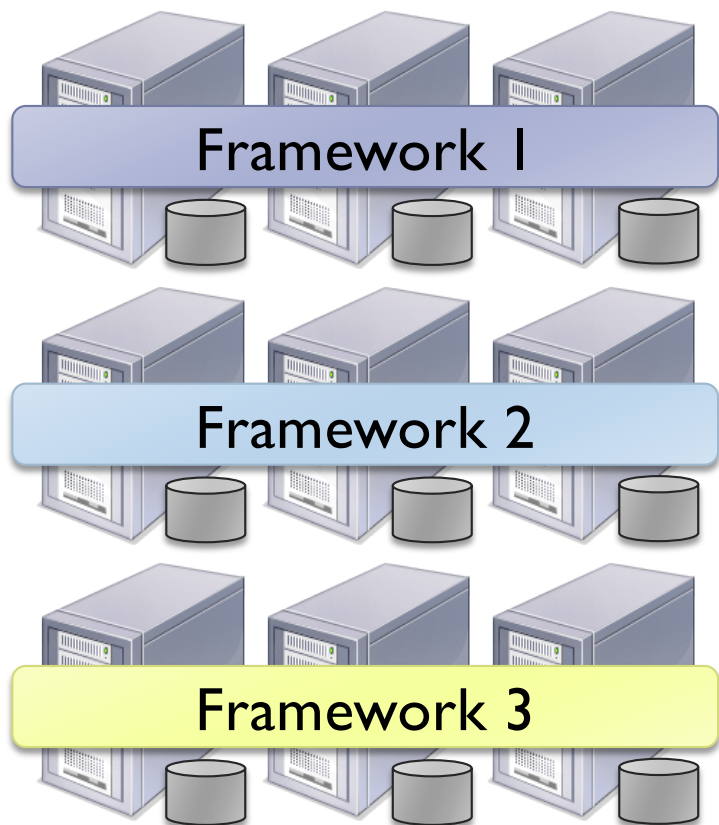
Resulting design: Small microkernel-like core that pushes scheduling logic to frameworks

Design Elements

- ▶ **Fine-grained sharing:**
 - ▶ Allocation at the level of tasks within a job
 - ▶ Improves utilization, latency, and data locality
- ▶ **Resource offers:**
 - ▶ Simple, scalable application-controlled scheduling mechanism

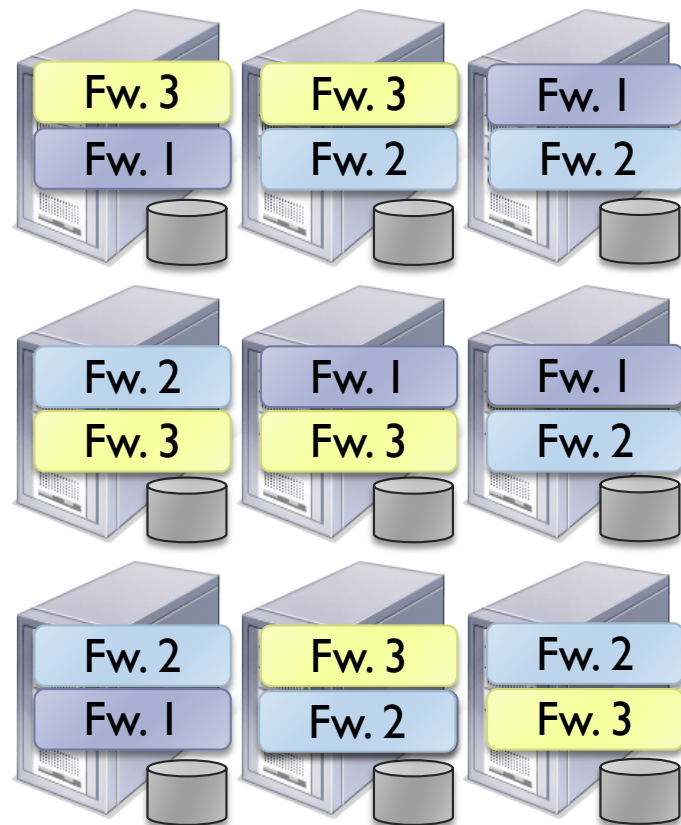
Element 1: Fine-Grained Sharing

Coarse-Grained Sharing (HPC):



Storage System (e.g. HDFS)

Fine-Grained Sharing (Mesos):



Storage System (e.g. HDFS)

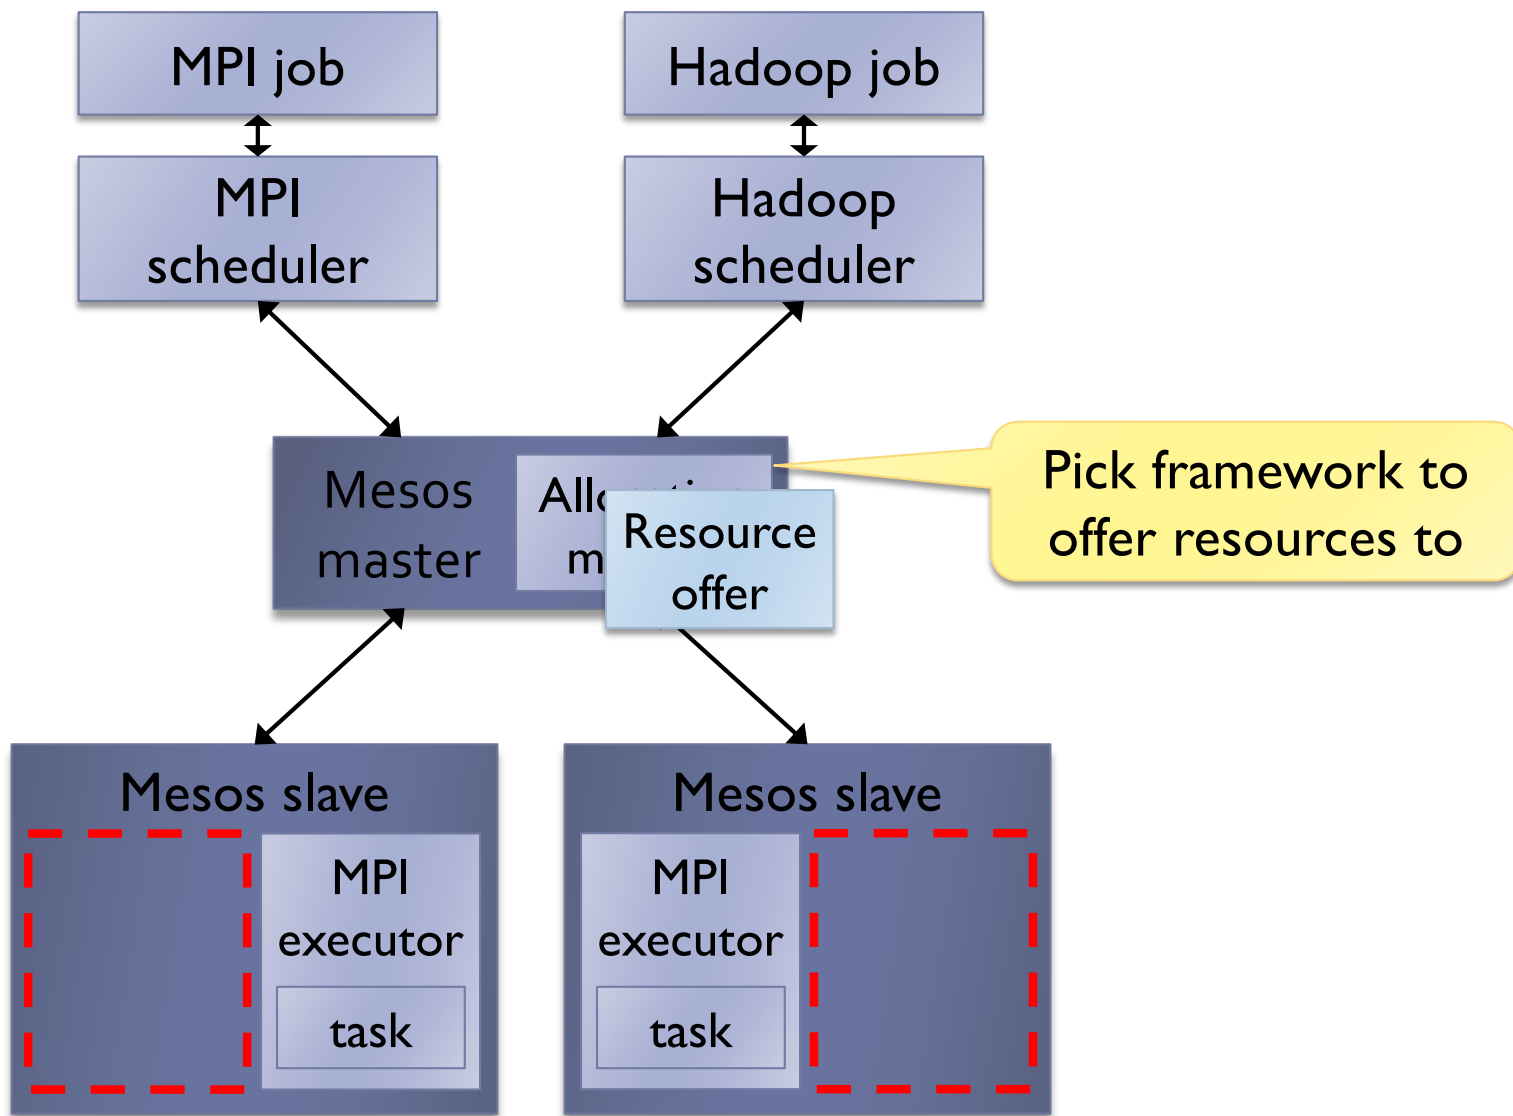
+ Improved utilization, responsiveness, data locality

MapReduce. Spark. Mesos. Yarn

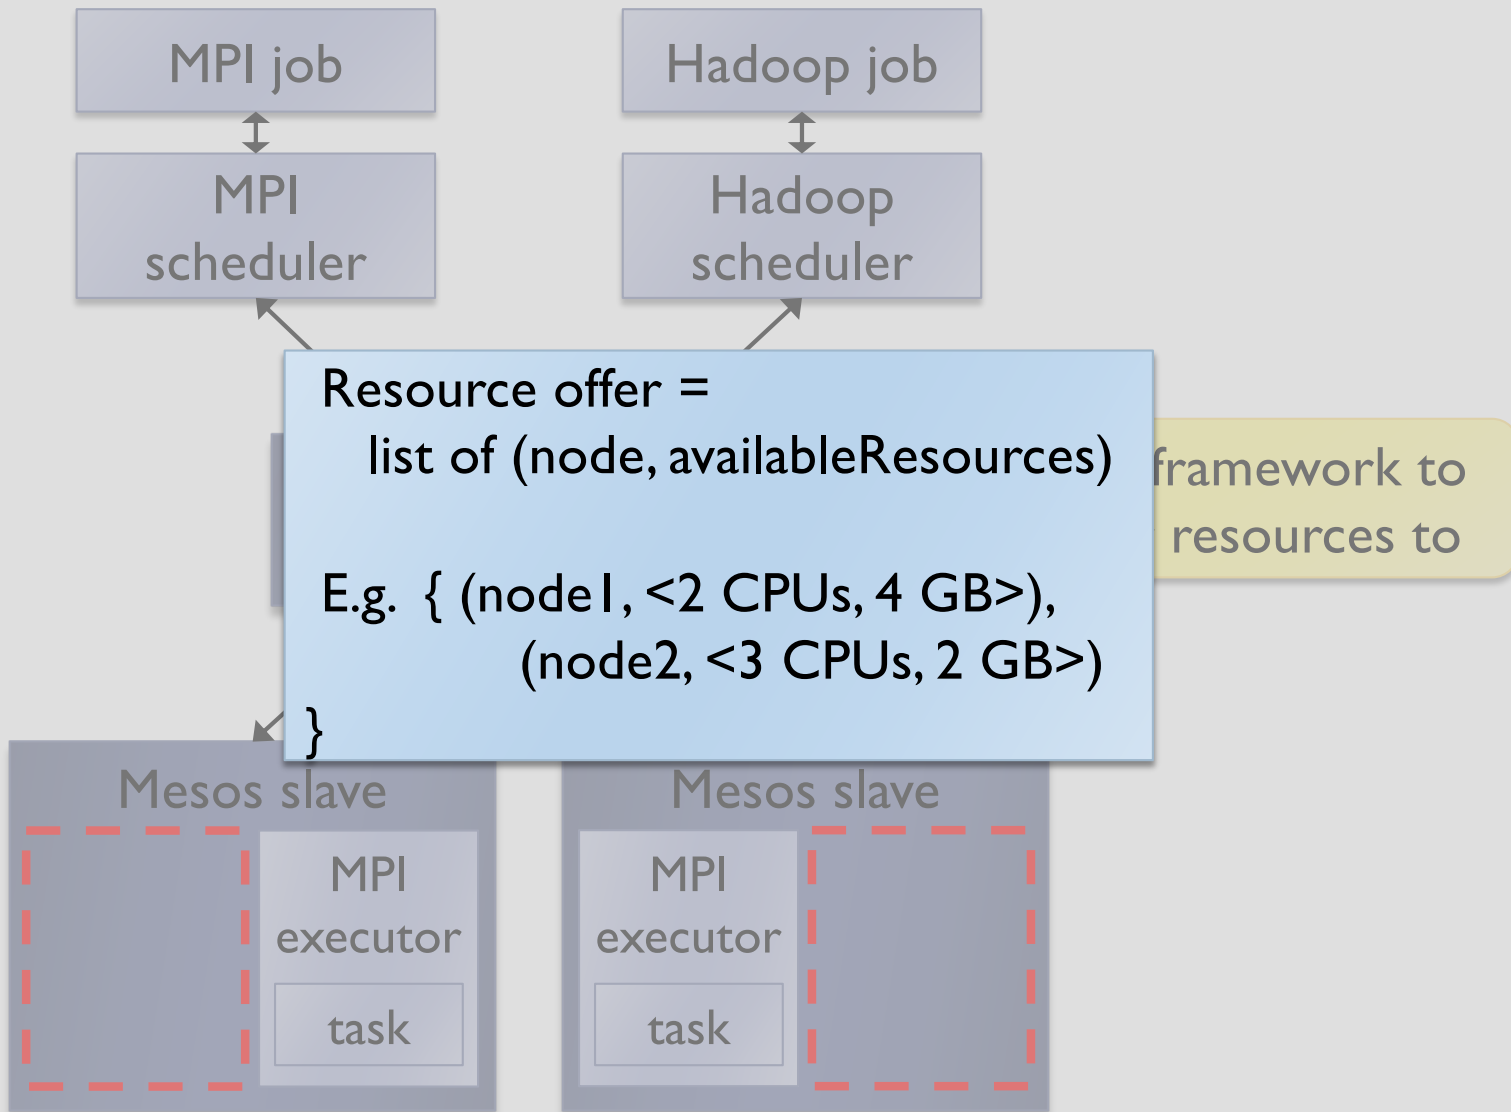
Element 2: Resource Offers

- ▶ **Option: Global scheduler**
 - ▶ Frameworks express needs in a specification language, global scheduler matches them to resources
 - ▶ + Can make optimal decisions
- ▶ **– Complex: language must support all framework needs**
 - ▶ – Difficult to scale and to make robust
 - ▶ – Future frameworks may have unanticipated needs

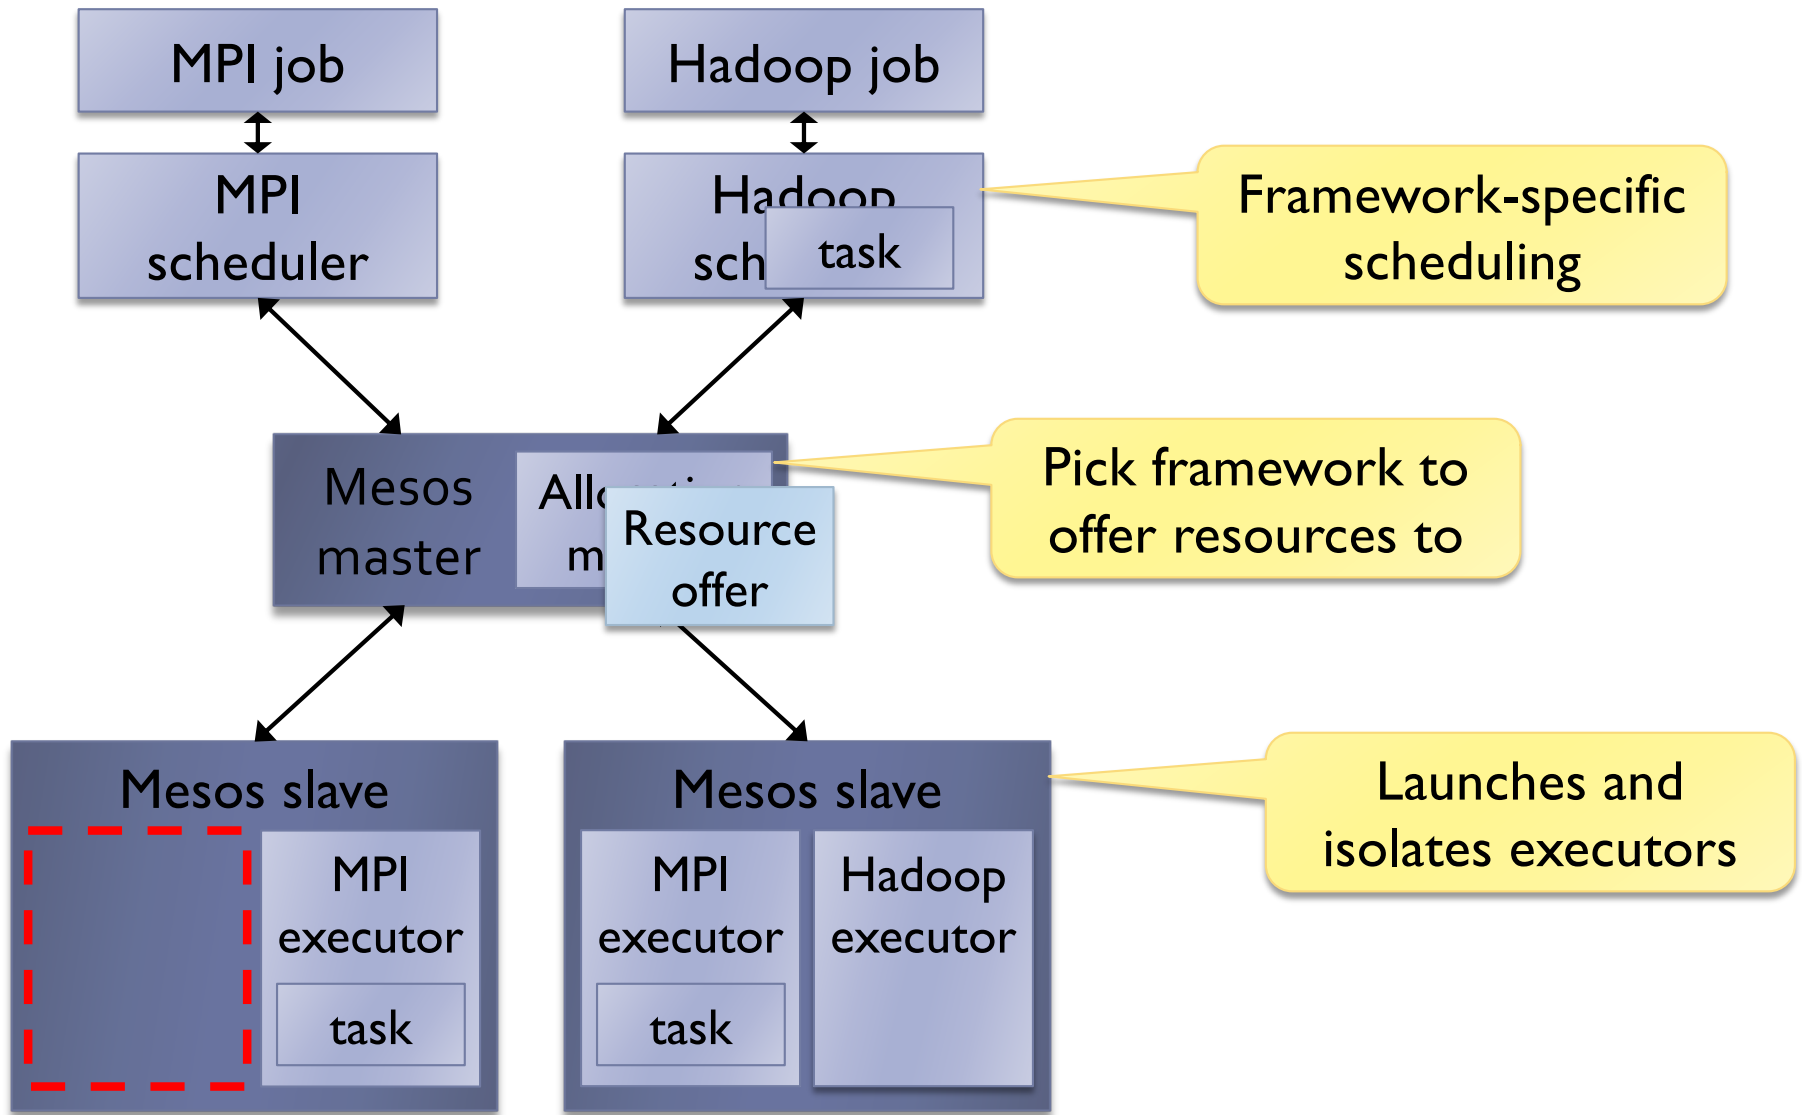
Mesos Architecture



Mesos Architecture



Mesos Architecture



Optimization: Filters

- ▶ Let frameworks short-circuit rejection by providing a predicate on resources to be offered
 - ▶ E.g. “nodes from list L” or “nodes with > 8 GB RAM”
 - ▶ Could generalize to other hints as well
- ▶ Ability to reject still ensures correctness when needs cannot be expressed using filters

Implementation Stats

- ▶ 20,000 lines of C++
- ▶ Master failover using ZooKeeper
- ▶ Frameworks ported: Hadoop, MPI, Torque
- ▶ New specialized framework: Spark, for iterative jobs (up to 20 × faster than Hadoop)

- ▶ Open source in Apache Incubator

Users

- ▶ Twitter uses Mesos on > 100 nodes to run ~ 12 production services (mostly stream processing)
- ▶ Berkeley machine learning researchers are running several algorithms at scale on Spark
- ▶ Conviva is using Spark for data analytics
- ▶ UCSF medical researchers are using Mesos to run Hadoop and eventually non-Hadoop apps

Framework Isolation

- ▶ Mesos uses OS isolation mechanisms, such as Linux containers and Solaris projects
- ▶ Containers currently support CPU, memory, IO and network bandwidth isolation
- ▶ Not perfect, but much better than no isolation

Analysis

- ▶ **Resource offers work well when:**
 - ▶ Frameworks can scale up and down elastically
 - ▶ Task durations are homogeneous
 - ▶ Frameworks have many preferred nodes
- ▶ **These conditions hold in current data analytics frameworks (MapReduce, Dryad, ...)**
 - ▶ Work divided into short tasks to facilitate load balancing and fault recovery
 - ▶ Data replicated across multiple nodes

Revocation

- ▶ Mesos allocation modules can revoke (kill) tasks to meet organizational SLOs
- ▶ Framework given a grace period to clean up
- ▶ “Guaranteed share” API lets frameworks avoid revocation by staying below a certain share

Mesos API

Scheduler Callbacks

resourceOffer(offerId, offers)
offerRescinded(offerId)
statusUpdate(taskId, status)
slaveLost(slaveId)

Scheduler Actions

replyToOffer(offerId, tasks)
setNeedsOffers(bool)
setFilters(filters)
getGuaranteedShare()
killTask(taskId)

Executor Callbacks

launchTask(taskDescriptor)
killTask(taskId)

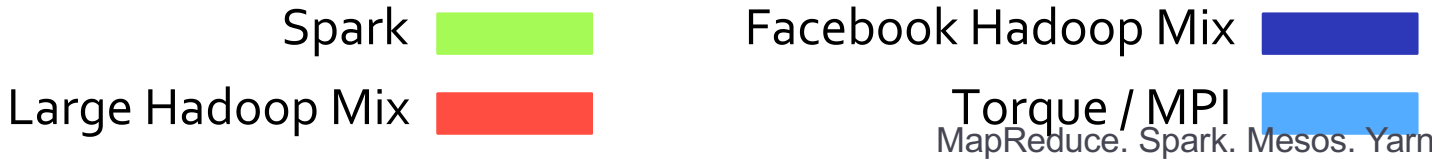
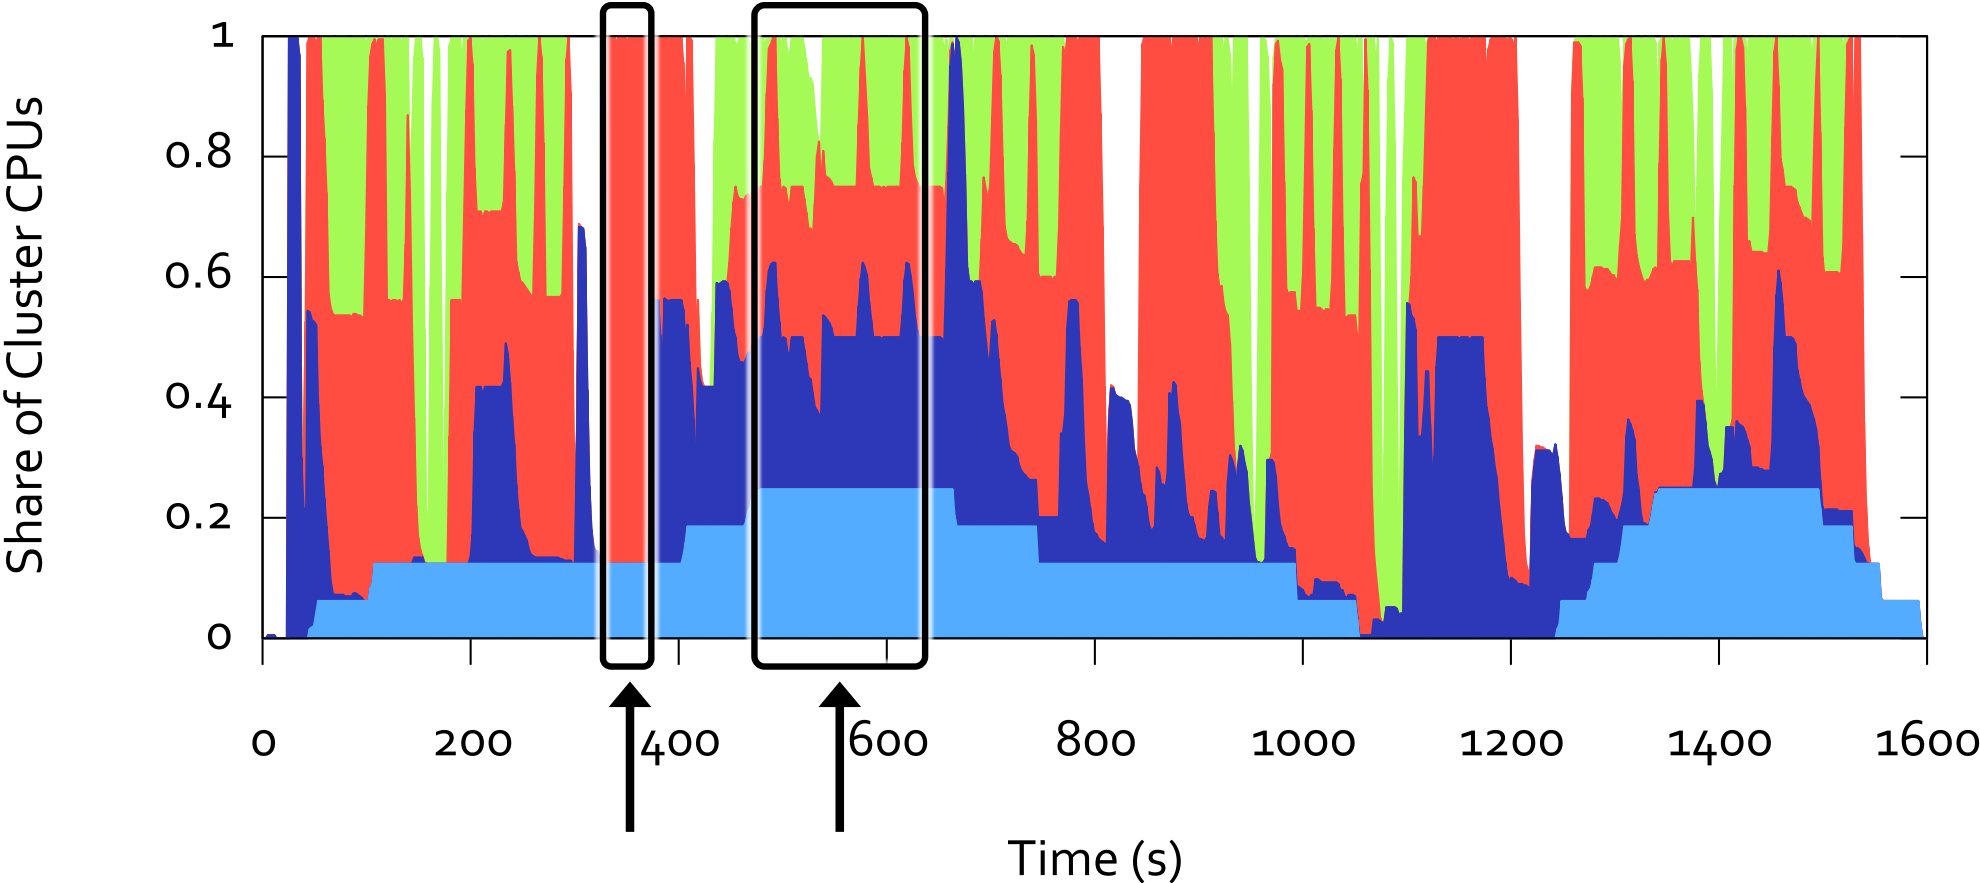
Executor Actions

sendStatus(taskId, status)

Evaluation

- ▶ Utilization and performance vs static partitioning
- ▶ Framework placement goals: data locality
- ▶ Scalability
- ▶ Fault recovery

Dynamic Resource Sharing



Mesos vs Static Partitioning

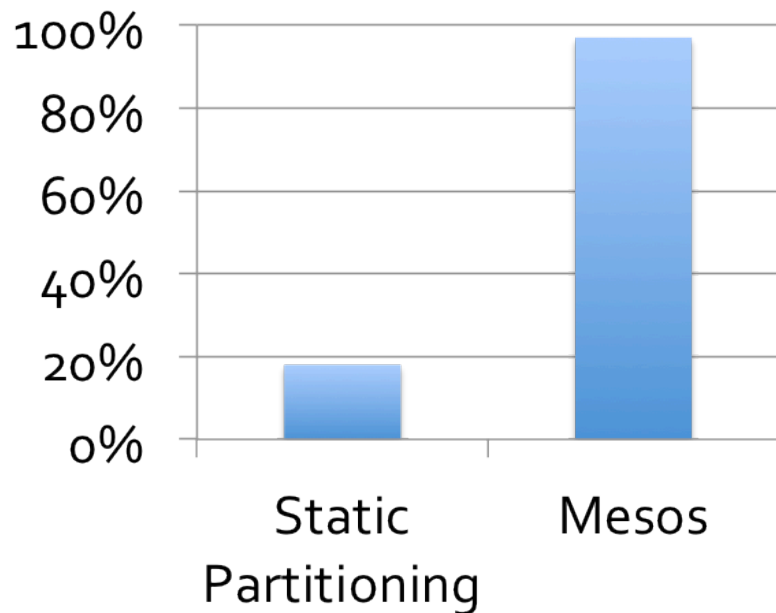
- ▶ Compared performance with statically partitioned cluster where each framework gets 25% of nodes

Framework	Speedup on Mesos
Facebook Hadoop Mix	1.14×
Large Hadoop Mix	2.10×
Spark	1.26×
Torque / MPI	0.96×

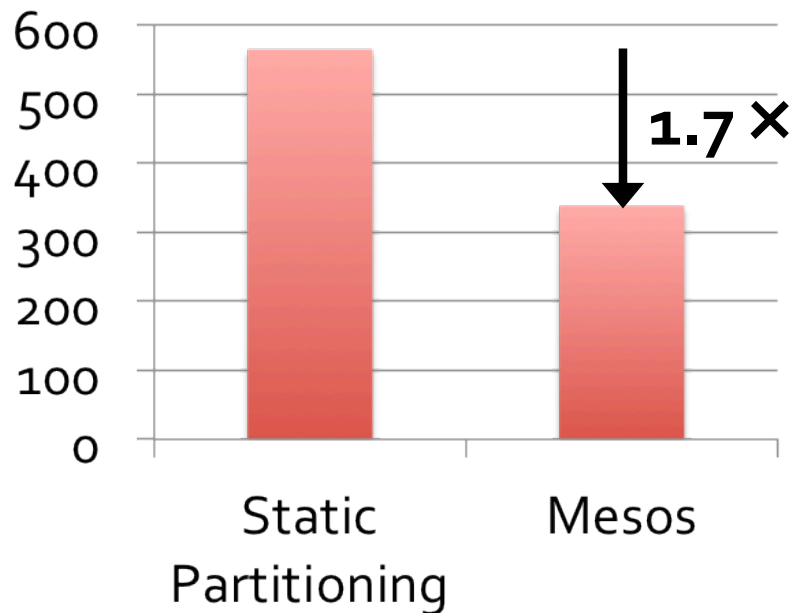
Data Locality with Resource Offers

- ▶ Ran 16 instances of Hadoop on a shared HDFS cluster
- ▶ Used delay scheduling [EuroSys '10] in Hadoop to get locality (wait a short time to acquire data-local nodes)

Local Map Tasks (%)



Job Duration (s)

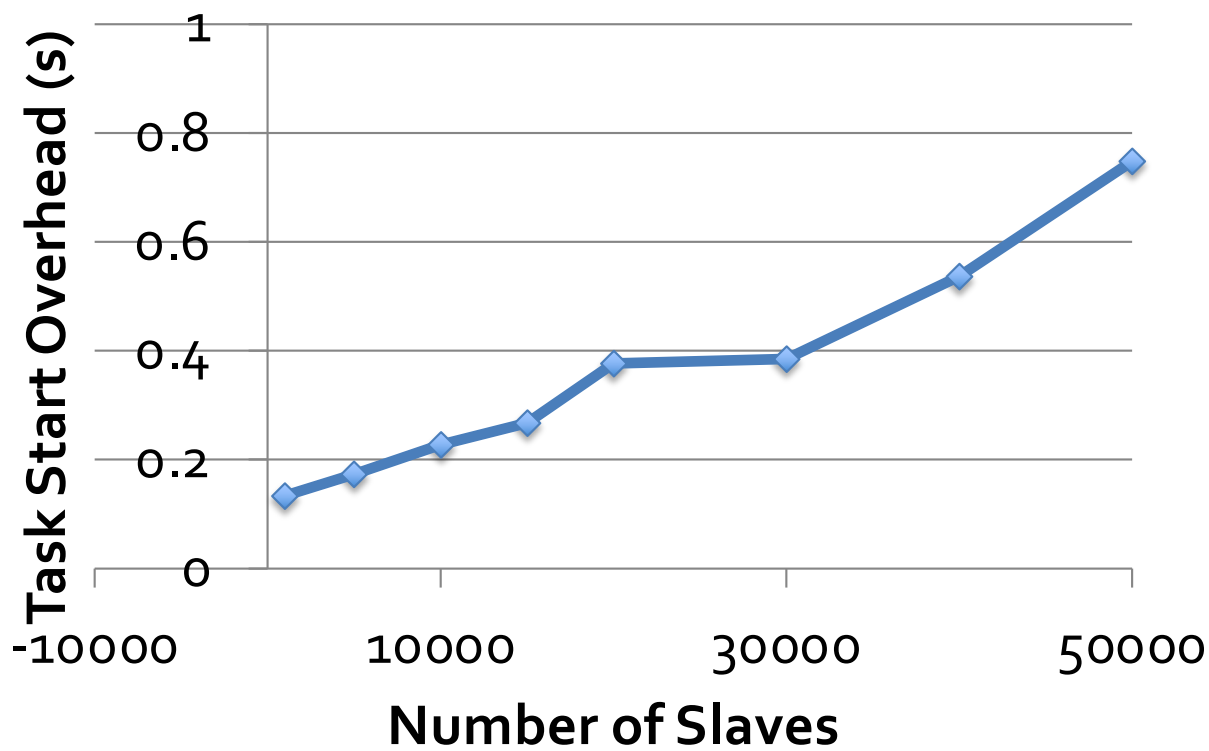


Scalability

- ▶ Mesos only performs inter-framework scheduling (e.g. fair sharing), which is easier than intra-framework scheduling

Result:

Scaled to 50,000 emulated slaves, 200 frameworks, 100K tasks (30s len)



Fault Tolerance

- ▶ Mesos master has only soft state: list of currently running frameworks and tasks
- ▶ Rebuild when frameworks and slaves re-register with new master after a failure
- ▶ Result: fault detection and recovery in ~10 sec

Conclusion

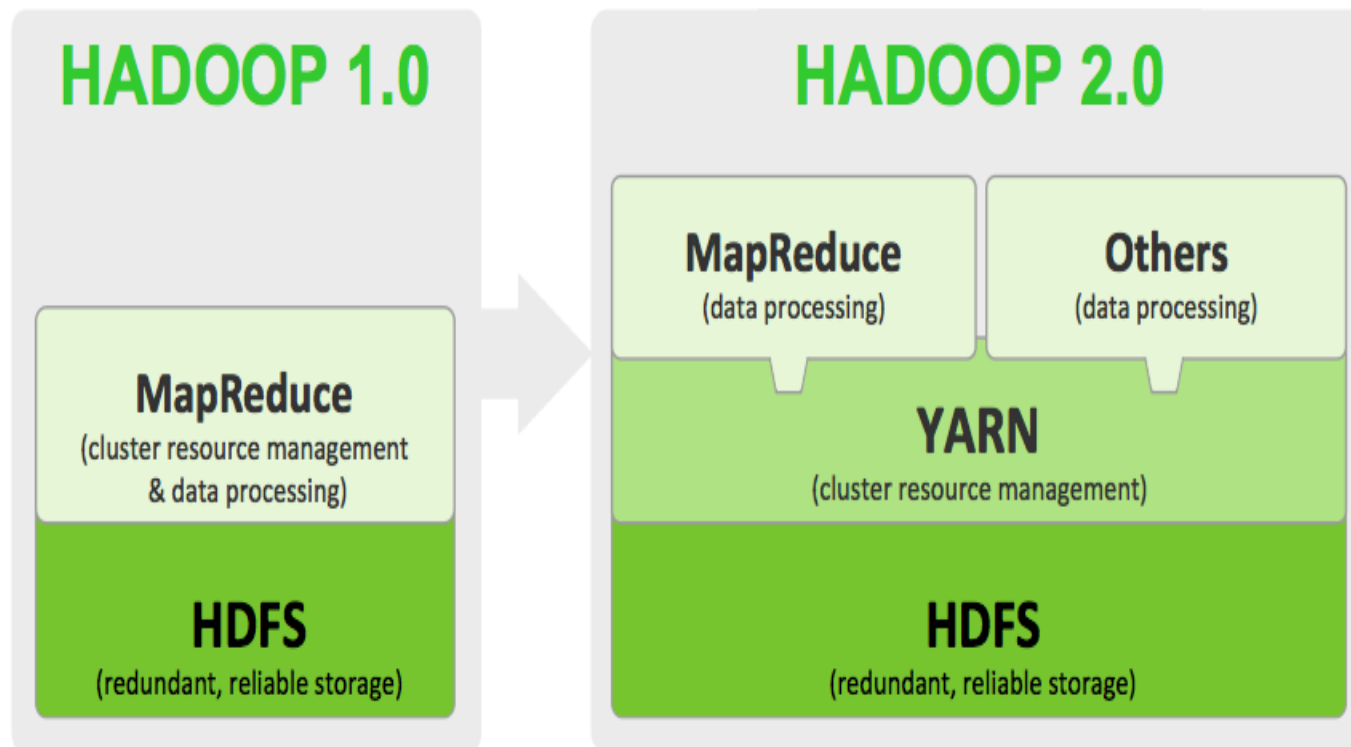
- ▶ **Mesos shares clusters efficiently among diverse frameworks thanks to two design elements:**
 - ▶ Fine-grained sharing at the level of tasks
 - ▶ Resource offers, a scalable mechanism for application-controlled scheduling
- ▶ **Enables co-existence of current frameworks and development of new specialized ones**
- ▶ **In use at Twitter, UC Berkeley, Conviva and UCSF**



4: Yarn

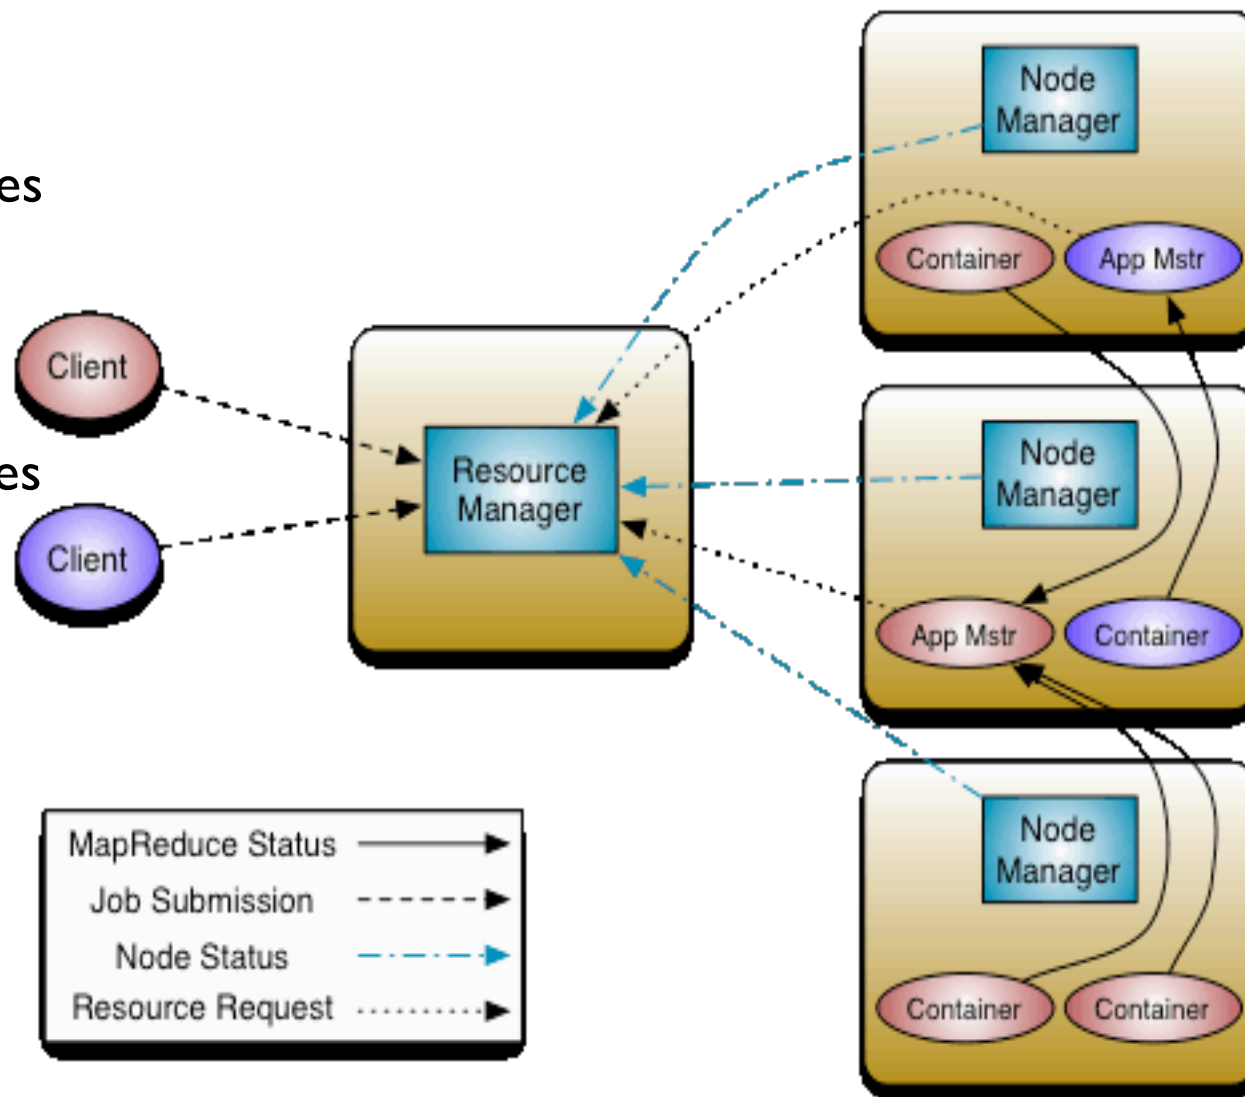
YARN - Yet Another Resource Negotiator

- ▶ Next version of MapReduce or MapReduce 2.0 (MRv2)
- ▶ In 2010 group at Yahoo! Began to design the next generation of MR



YARN architecture

- **Resource Manager**
 - Central Agent – Manages and allocates cluster resources
- **Node Manager**
 - Per-node agent – Manages and enforces node resource allocations
- **Application Master**
 - Per Application
 - Manages application life cycle and task scheduling



YARN – Resource Manager Failure

- ▶ After a crash a new Resource Manager instance needs to be brought up (by an administrator)
- ▶ It recovers from saved state
- ▶ State consists of
 - ▶ node managers in the systems
 - ▶ running applications
- ▶ State to manage is much more manageable than that of Job Tracker.
 - ▶ Tasks are not part of Resource Managers state.
 - ▶ They are handled by the application master.