

COMP 322: Parallel and Concurrent Programming

Lecture 6: Map/Reduce

Mack Joyner
mjoyner@rice.edu

<http://comp322.rice.edu>



Streaming data requirements have skyrocketed

AT&T processes roughly 30 petabytes per day through its telecommunications network

Facebook, Amazon, Twitter, etc, have comparable throughputs

IBM Watson knowledge base stored roughly 4 terabytes of data when winning at Jeopardy

The amount of data in the world was estimated to be 44 zettabytes at the beginning of 2020.

175 zettabytes by 2025.

By 2025, the amount of data generated each day is expected to reach 463 exabytes globally.

Electronic Arts process roughly 50 terabytes of data every day.

In 2019, nearly 695,000 hours of Netflix content were watched per minute across the world.




Parallelism enables processing of big data

Continuously streaming data needs to be processed at least as fast as it is accumulated, or we will never catch up

The bottleneck in processing very large data sets is dominated by the speed of disk access

More processors accessing more disks enables faster processing



But massive parallel programming is hard!

100's of 1000's of servers

Dozens of cores per server

Gigabytes (terabytes) of data per server

Where to begin?

“MapReduce: Simplified Data Processing on Large Clusters”. *Jeffrey Dean and Sanjay Ghemawat.*

OSDI'04: Sixth Symposium on Operating System Design and Implementation, San Francisco,

CA (2004), pp. 137-150

"Our abstraction is inspired by the map and reduce primitives present in Lisp and many other functional languages."



Map/Reduce frameworks

Hadoop MapReduce

Java based

All data is stored on Hadoops distributed file system (HDFS)

Can process enormous data sets

More in COMP 330!

Apache Spark

Fits operations in memory

APIs for Java, Scala, Python and R

Requires a lot more computer memory

More in COMP 330!



MapReduce Pattern

Apply **Map** function f to user supplied record of key-value pairs

Compute set of intermediate key/value pairs

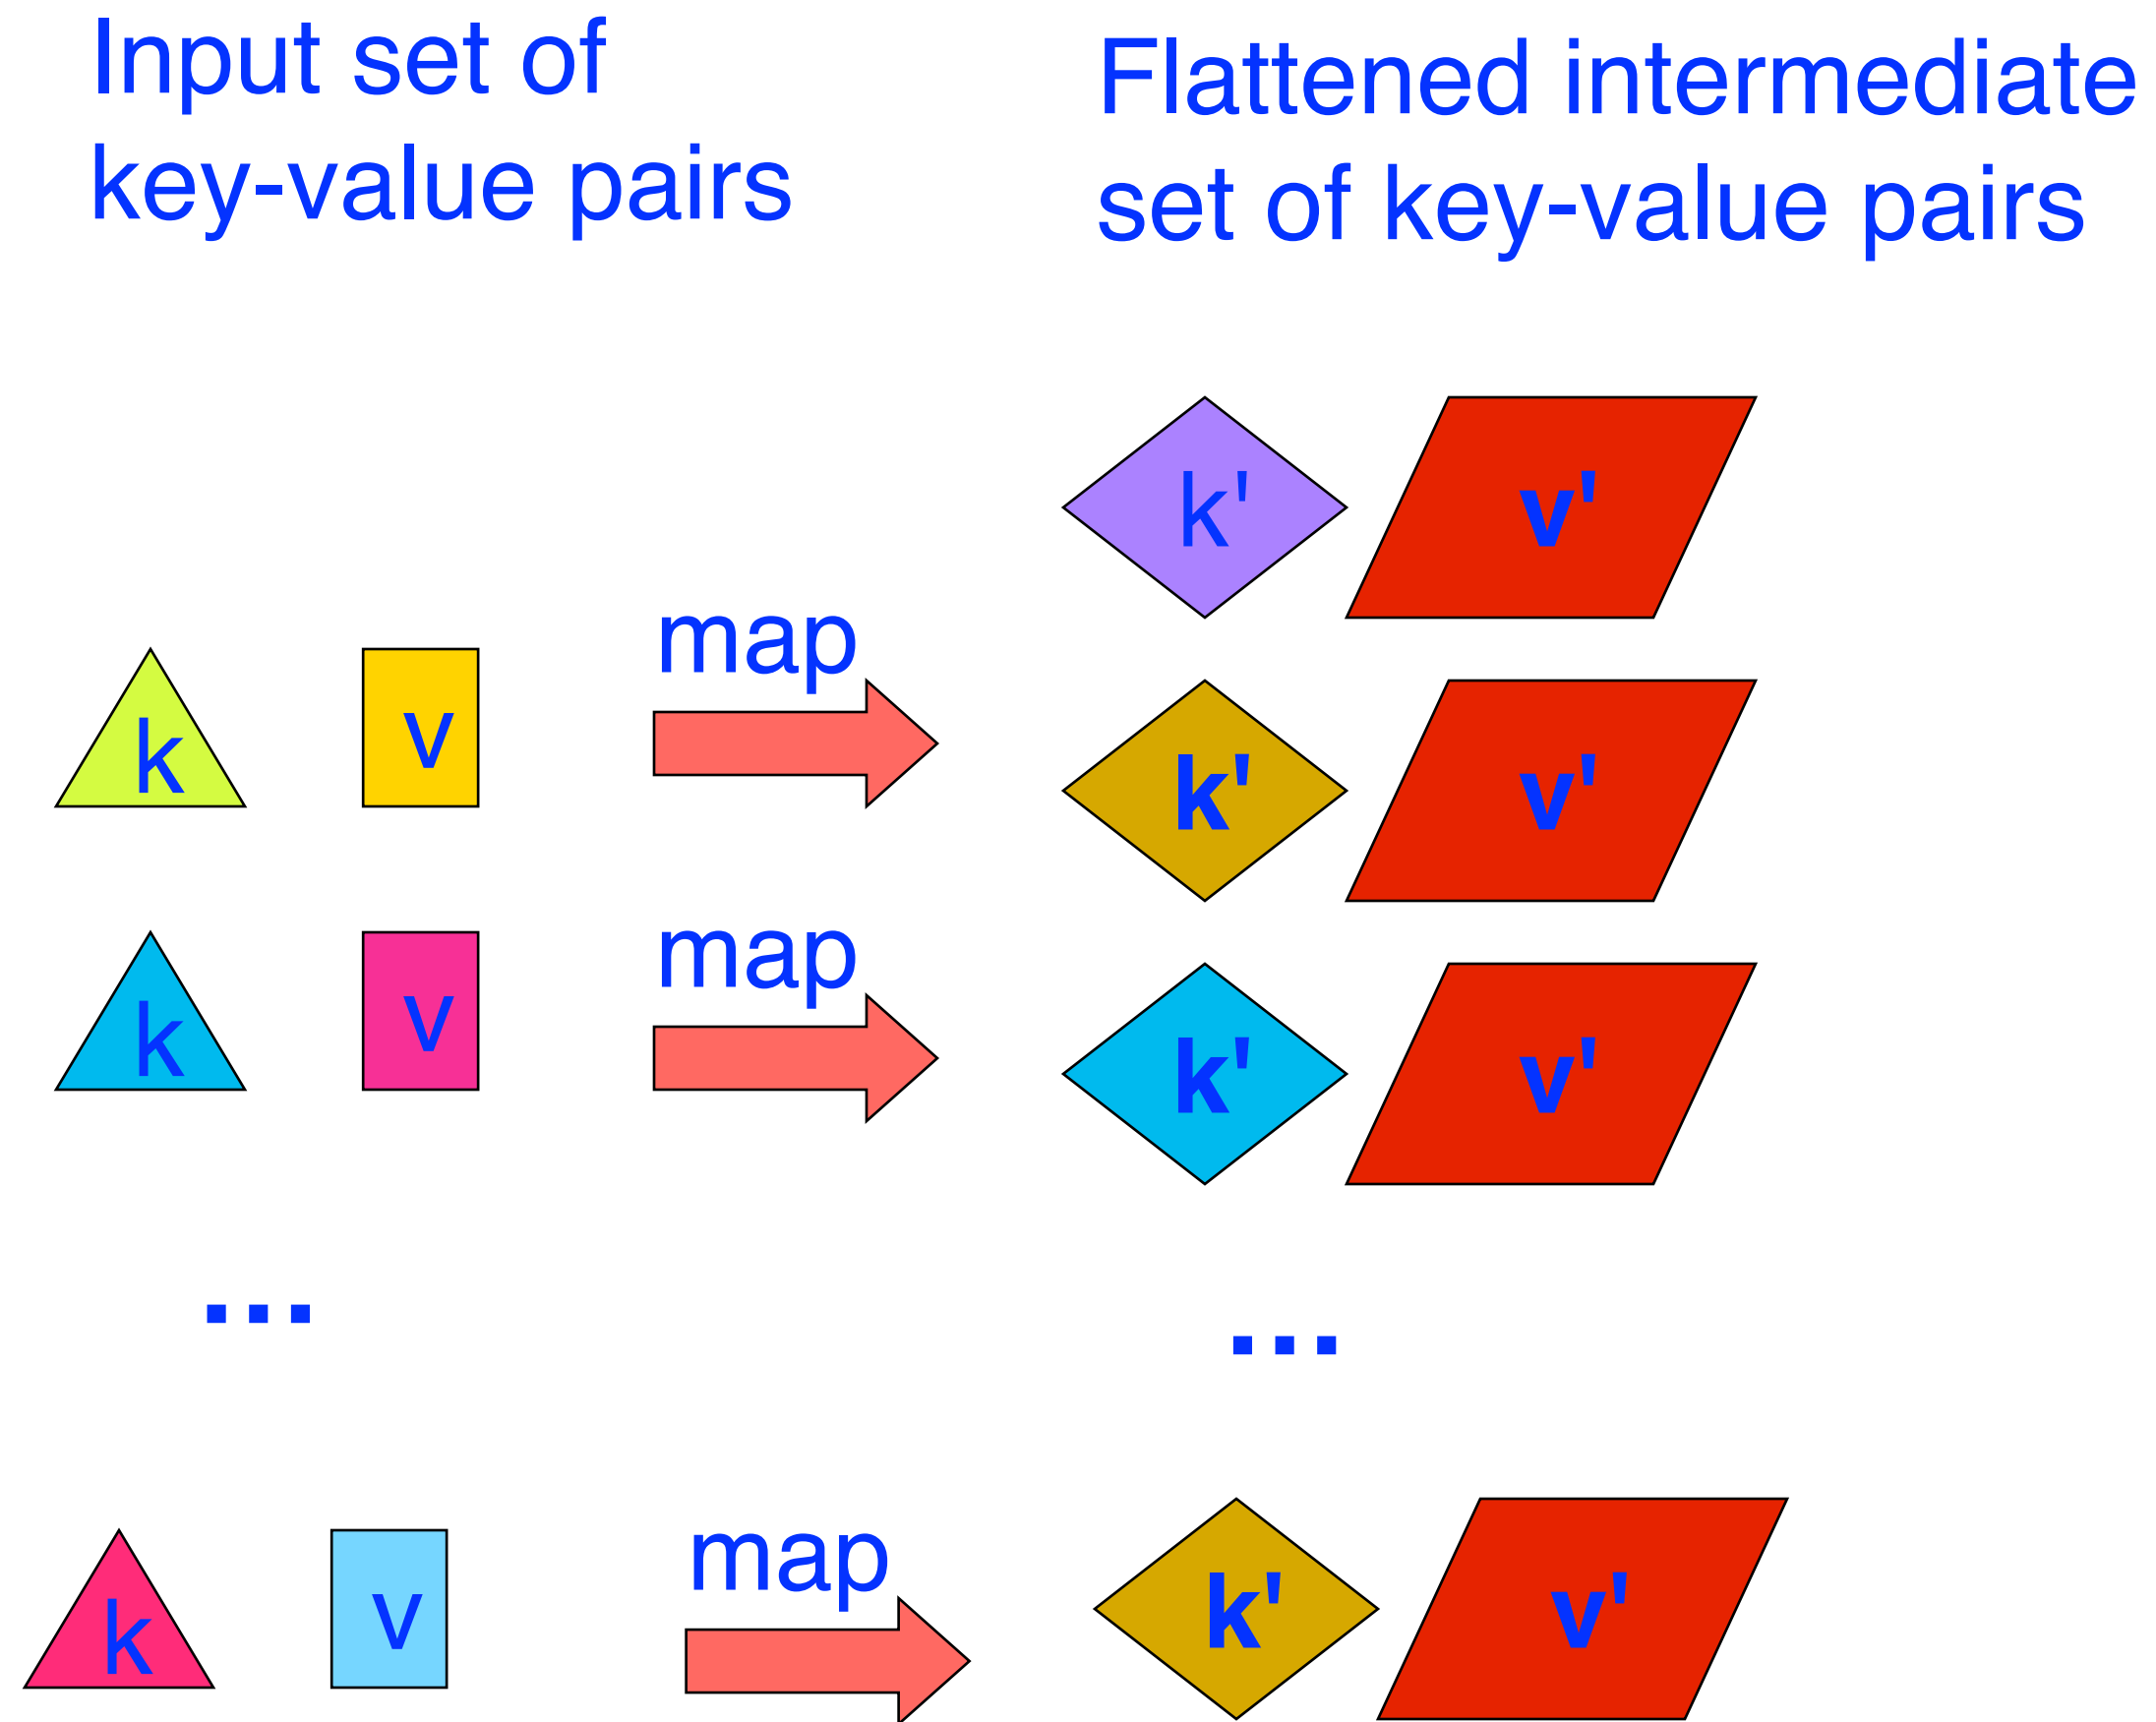
Apply **Reduce** operation g to all values that share the same key to combine derived data properly

Often produces smaller set of values

User supplies Map and Reduce operations in a functional model so that the system can parallelize them, and also re-execute them for fault tolerance



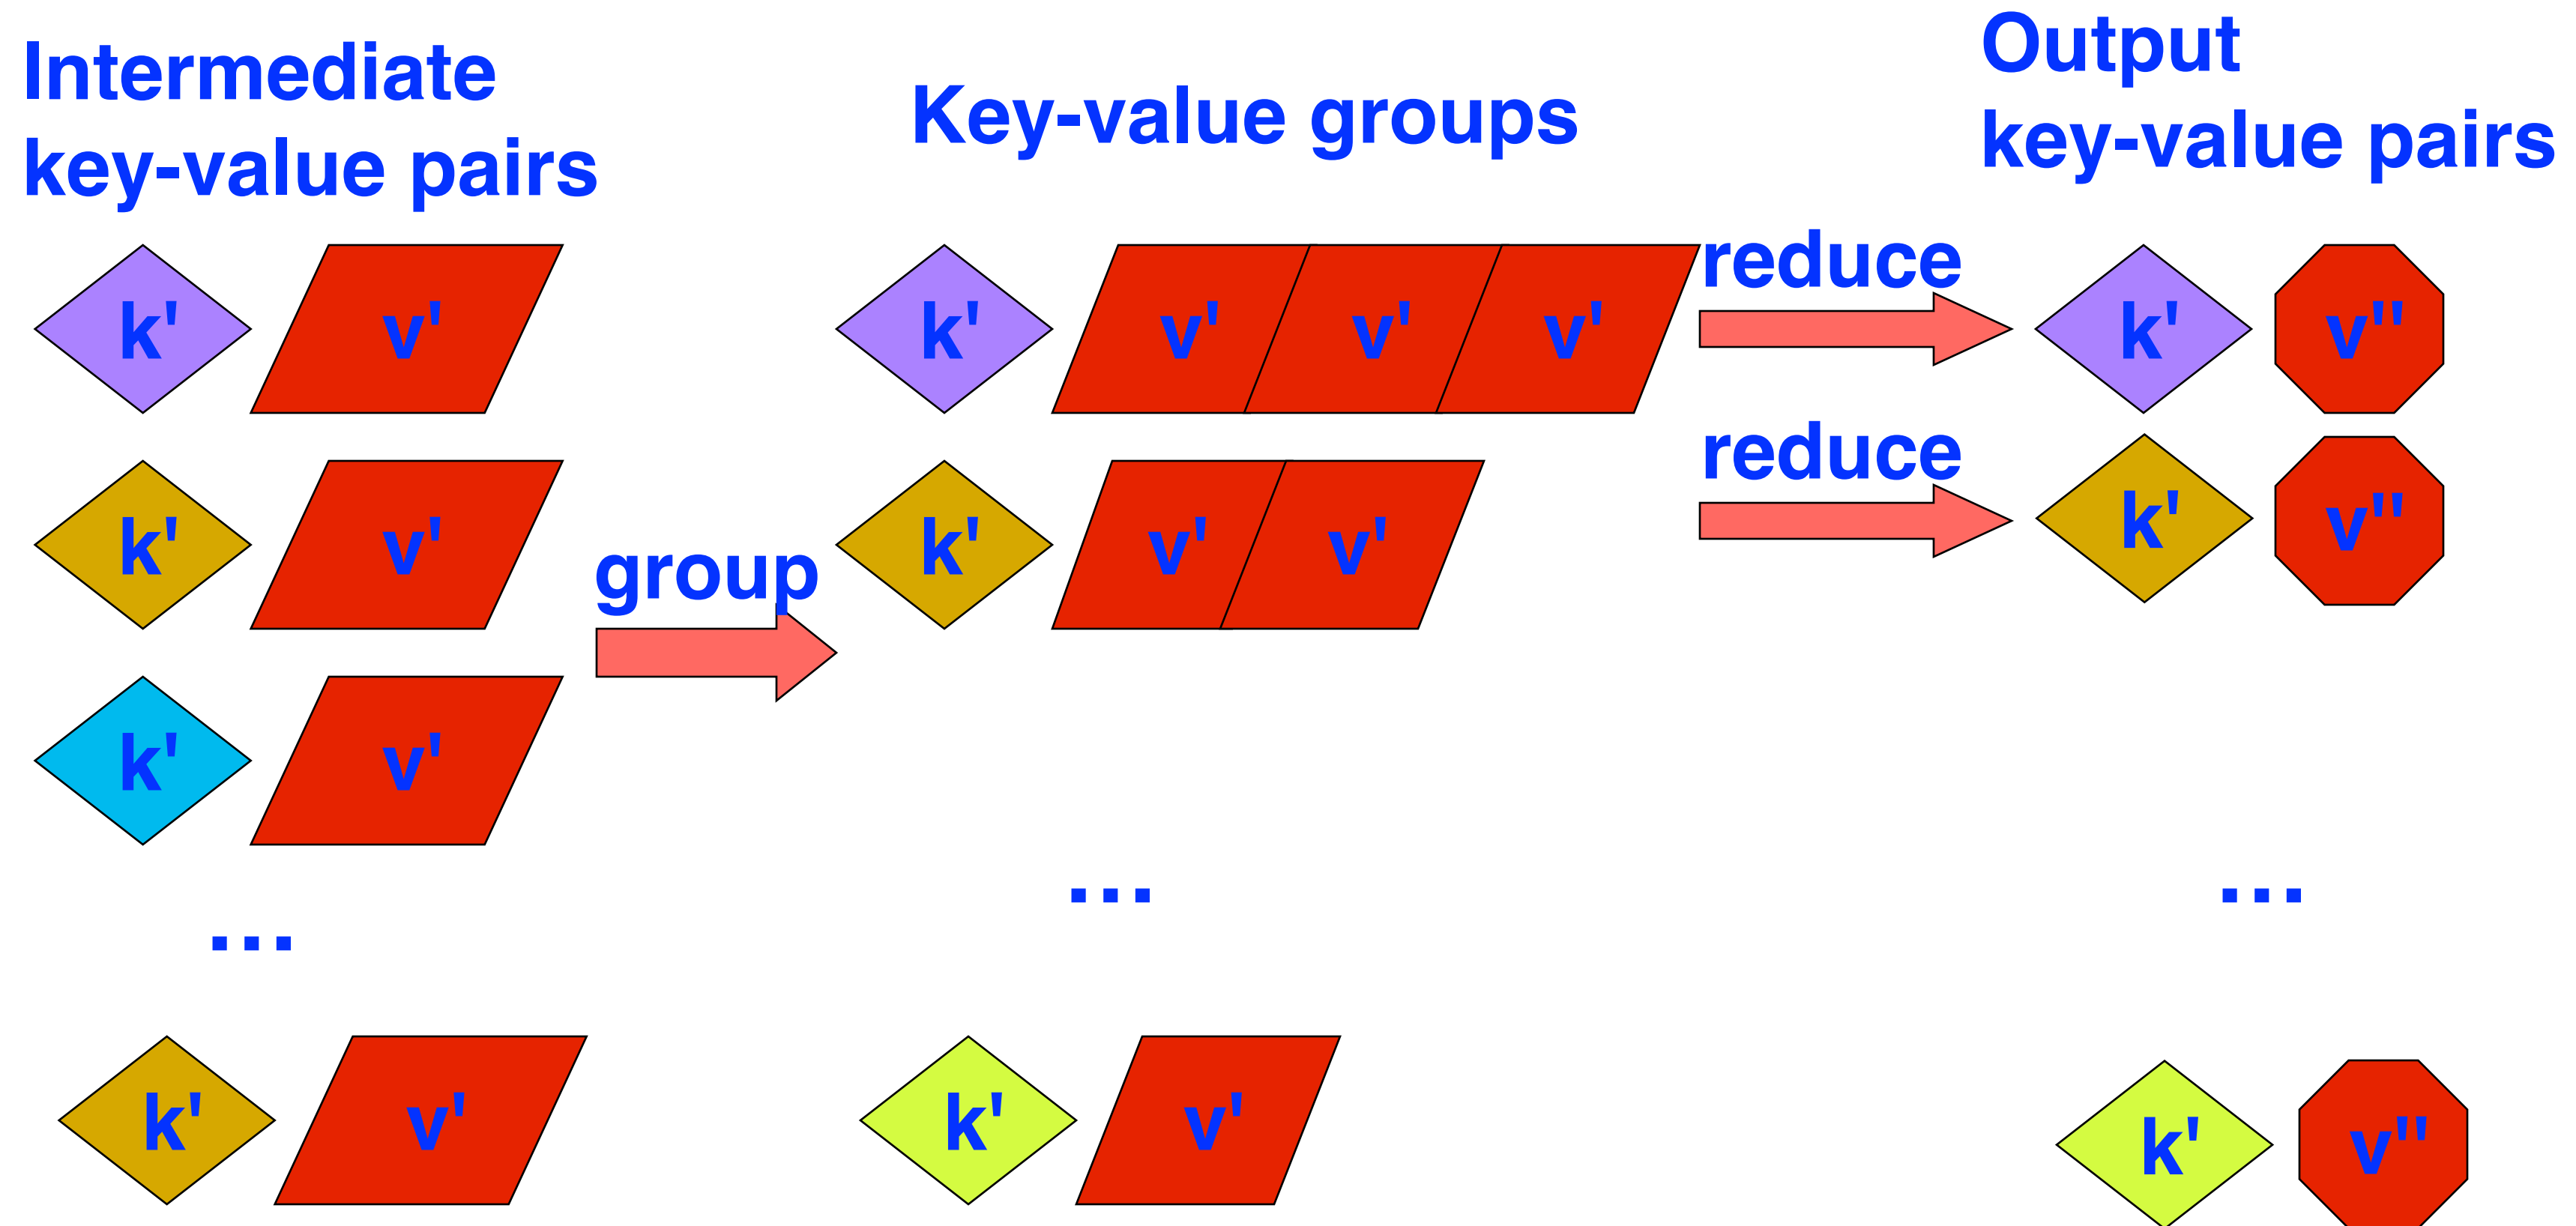
MapReduce: Map Step



Source: <http://infolab.stanford.edu/~ullman/mining/2009/mapreduce.ppt>



MapReduce: Reduce Step



Source: <http://infolab.stanford.edu/~ullman/mining/2009/mapreduce.ppt>



Map Reduce: Summary

Input set is of the form $\{(k_1, v_1), \dots, (k_n, v_n)\}$, where (k_i, v_i) consists of a key, k_i , and a value, v_i .

Assume key and value objects are immutable

Map function f generates sets of intermediate key-value pairs, $f(k_i, v_i) = \{(k_1', v_1'), \dots, (k_{m'}, v_{m'})\}$. The $k_{m'}$ keys can be different from k_i key in the map function.

Assume that a flatten operation is performed as a post-pass after the map operations, so as to avoid dealing with a set of sets. In other words, assume that a FlatMap operation is used.

Reduce operation groups together intermediate key-value pairs, $\{(k', v_j')\}$ with the same k' , and generates a reduced key-value pair, (k', v'') , for each such k' , using reduce function g



Google Uses MapReduce

Web crawl: Find outgoing links from HTML documents, aggregate by target document

Google Earth: Stitching overlapping satellite images to remove seams and to select high-quality imagery

Google Maps: Processing all road segments on Earth and render map tile images that display segments

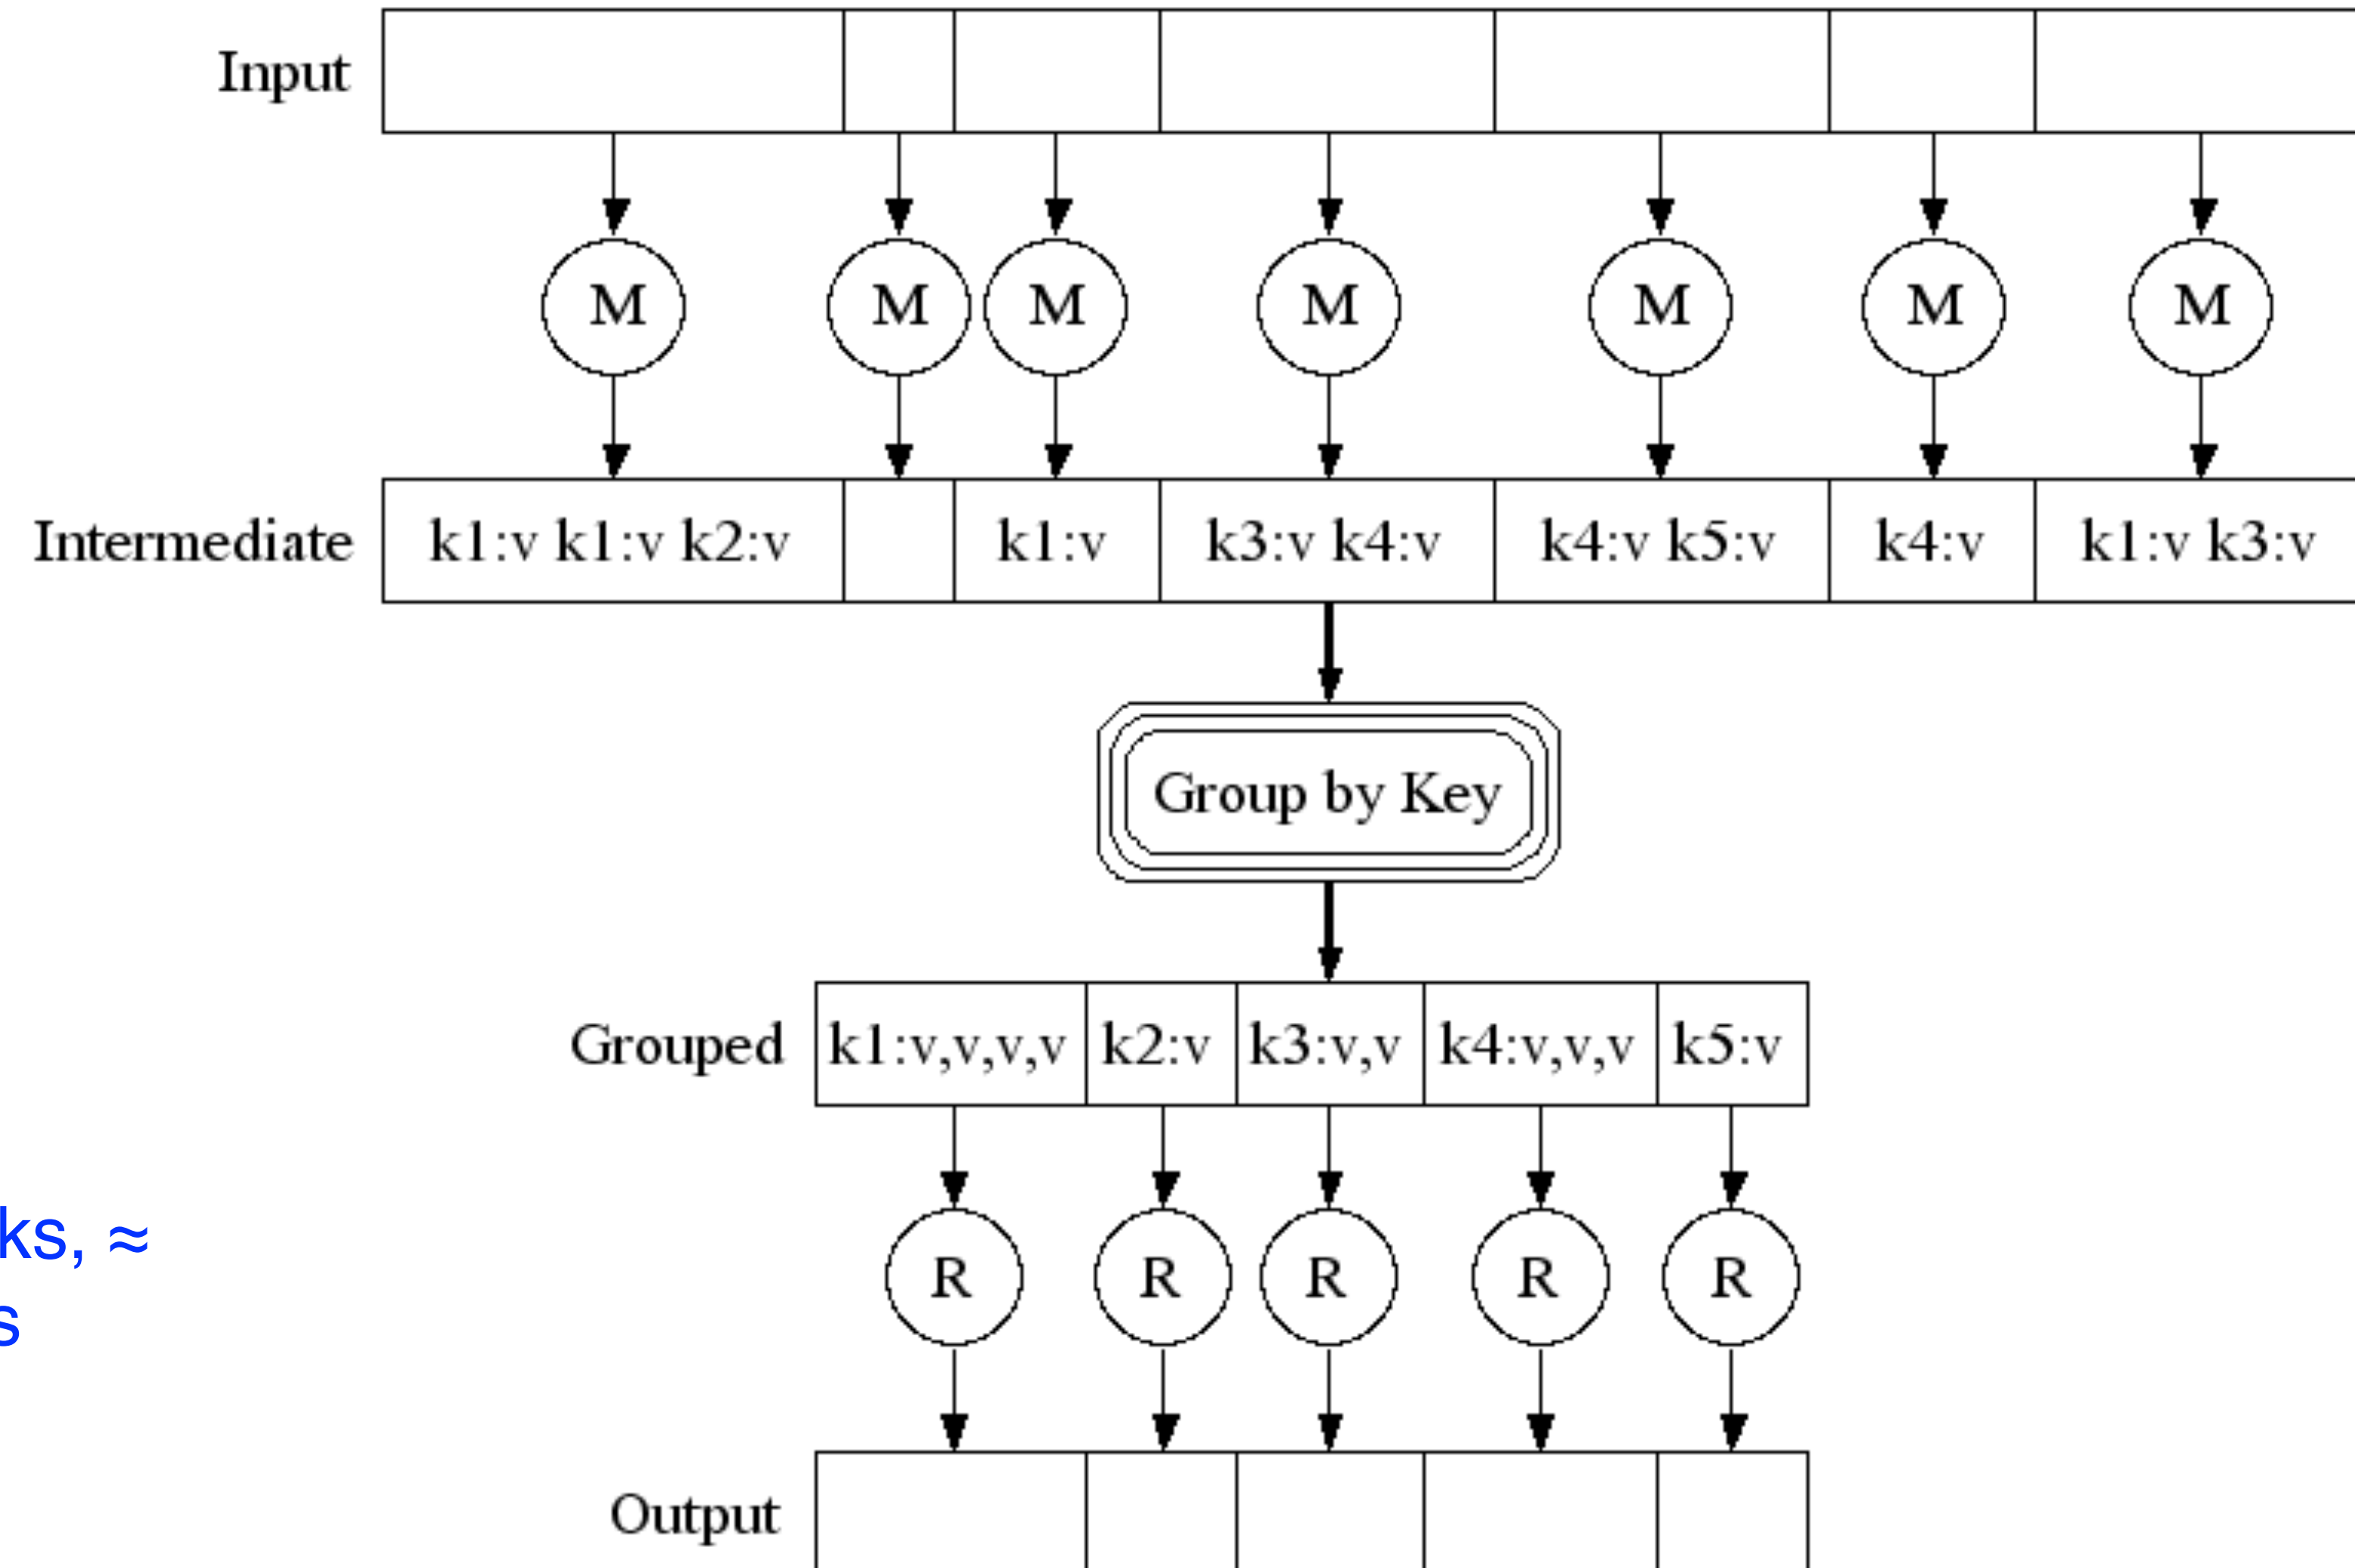


MapReduce Execution

Fine granularity tasks: many more map tasks than machines

Bucket sort to get same keys together

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



Word Count Example

In: set of words

Out: set of (word,count) pairs

Algorithm:

1. For each in word W , emit $(W, 1)$ as a key-value pair (map step).
2. Group together all key-value pairs with the same key (reduce step).
3. Perform a sum reduction on all values with the same key(reduce step).

All map operations in step 1 can execute in parallel with only local data accesses

Step 2 may involve a major reshuffle of data as all key-value pairs with the same key are grouped together.

Step 3 performs a standard reduction algorithm for all values with the same key, and in parallel for different keys.




Pseudocode for Word Count

1. `<String, Integer> map(String inKey, String inValue):`
2. `// inKey: document name`
3. `// inValue: document contents`
4. `for each word w in inValue:`
5. `emitIntermediate(w, 1) // Produce count of words`
- 6.
7. `<Integer> reduce(String outKey, Iterator<Integer> values):`
8. `// outKey: a word`
9. `// values: a list of counts`
10. `Integer result = 0`
11. `for each v in values:`
12. `result += v // the value from map was an integer`
13. `emit(result)`



Example Execution of Word Count Program

Distribute



Shuffle



Collect

is 6; it 2; not 2; that 5



Simple parallel WordCount using Streams

```
var file = Arrays.asList("that", "that", "is",  
    "is", "that", "that", "is", "not", "is", "not", "is", "that", "it", "it", "is");
```

```
Map<String, Integer> result = file.parallelStream()  
    .collect(groupingBy(Function.identity(), summingInt(e -> 1)));
```

```
result.forEach((k, v) -> System.out.println("The word \" + k + "\" appears " + v + " times"));
```

Shuffle

Reduce

Map

The word "that" appears 5 times
The word "not" appears 2 times
The word "is" appears 6 times
The word "it" appears 2 times



Summary

Map/Reduce is a programming model

Limited expressiveness

If you can solve your problem using the model, and your problem is very large, go for it

Heavily used in industry for processing large data

Hadoop, Spark, many others...

Heavily based on functional programming ideas

map, filter, fold, lambdas

Employs very similar ideas to Java Streams

Easy to parallelize across large number of servers

