


Map Reduce



Typical Application



What if...

INPUT

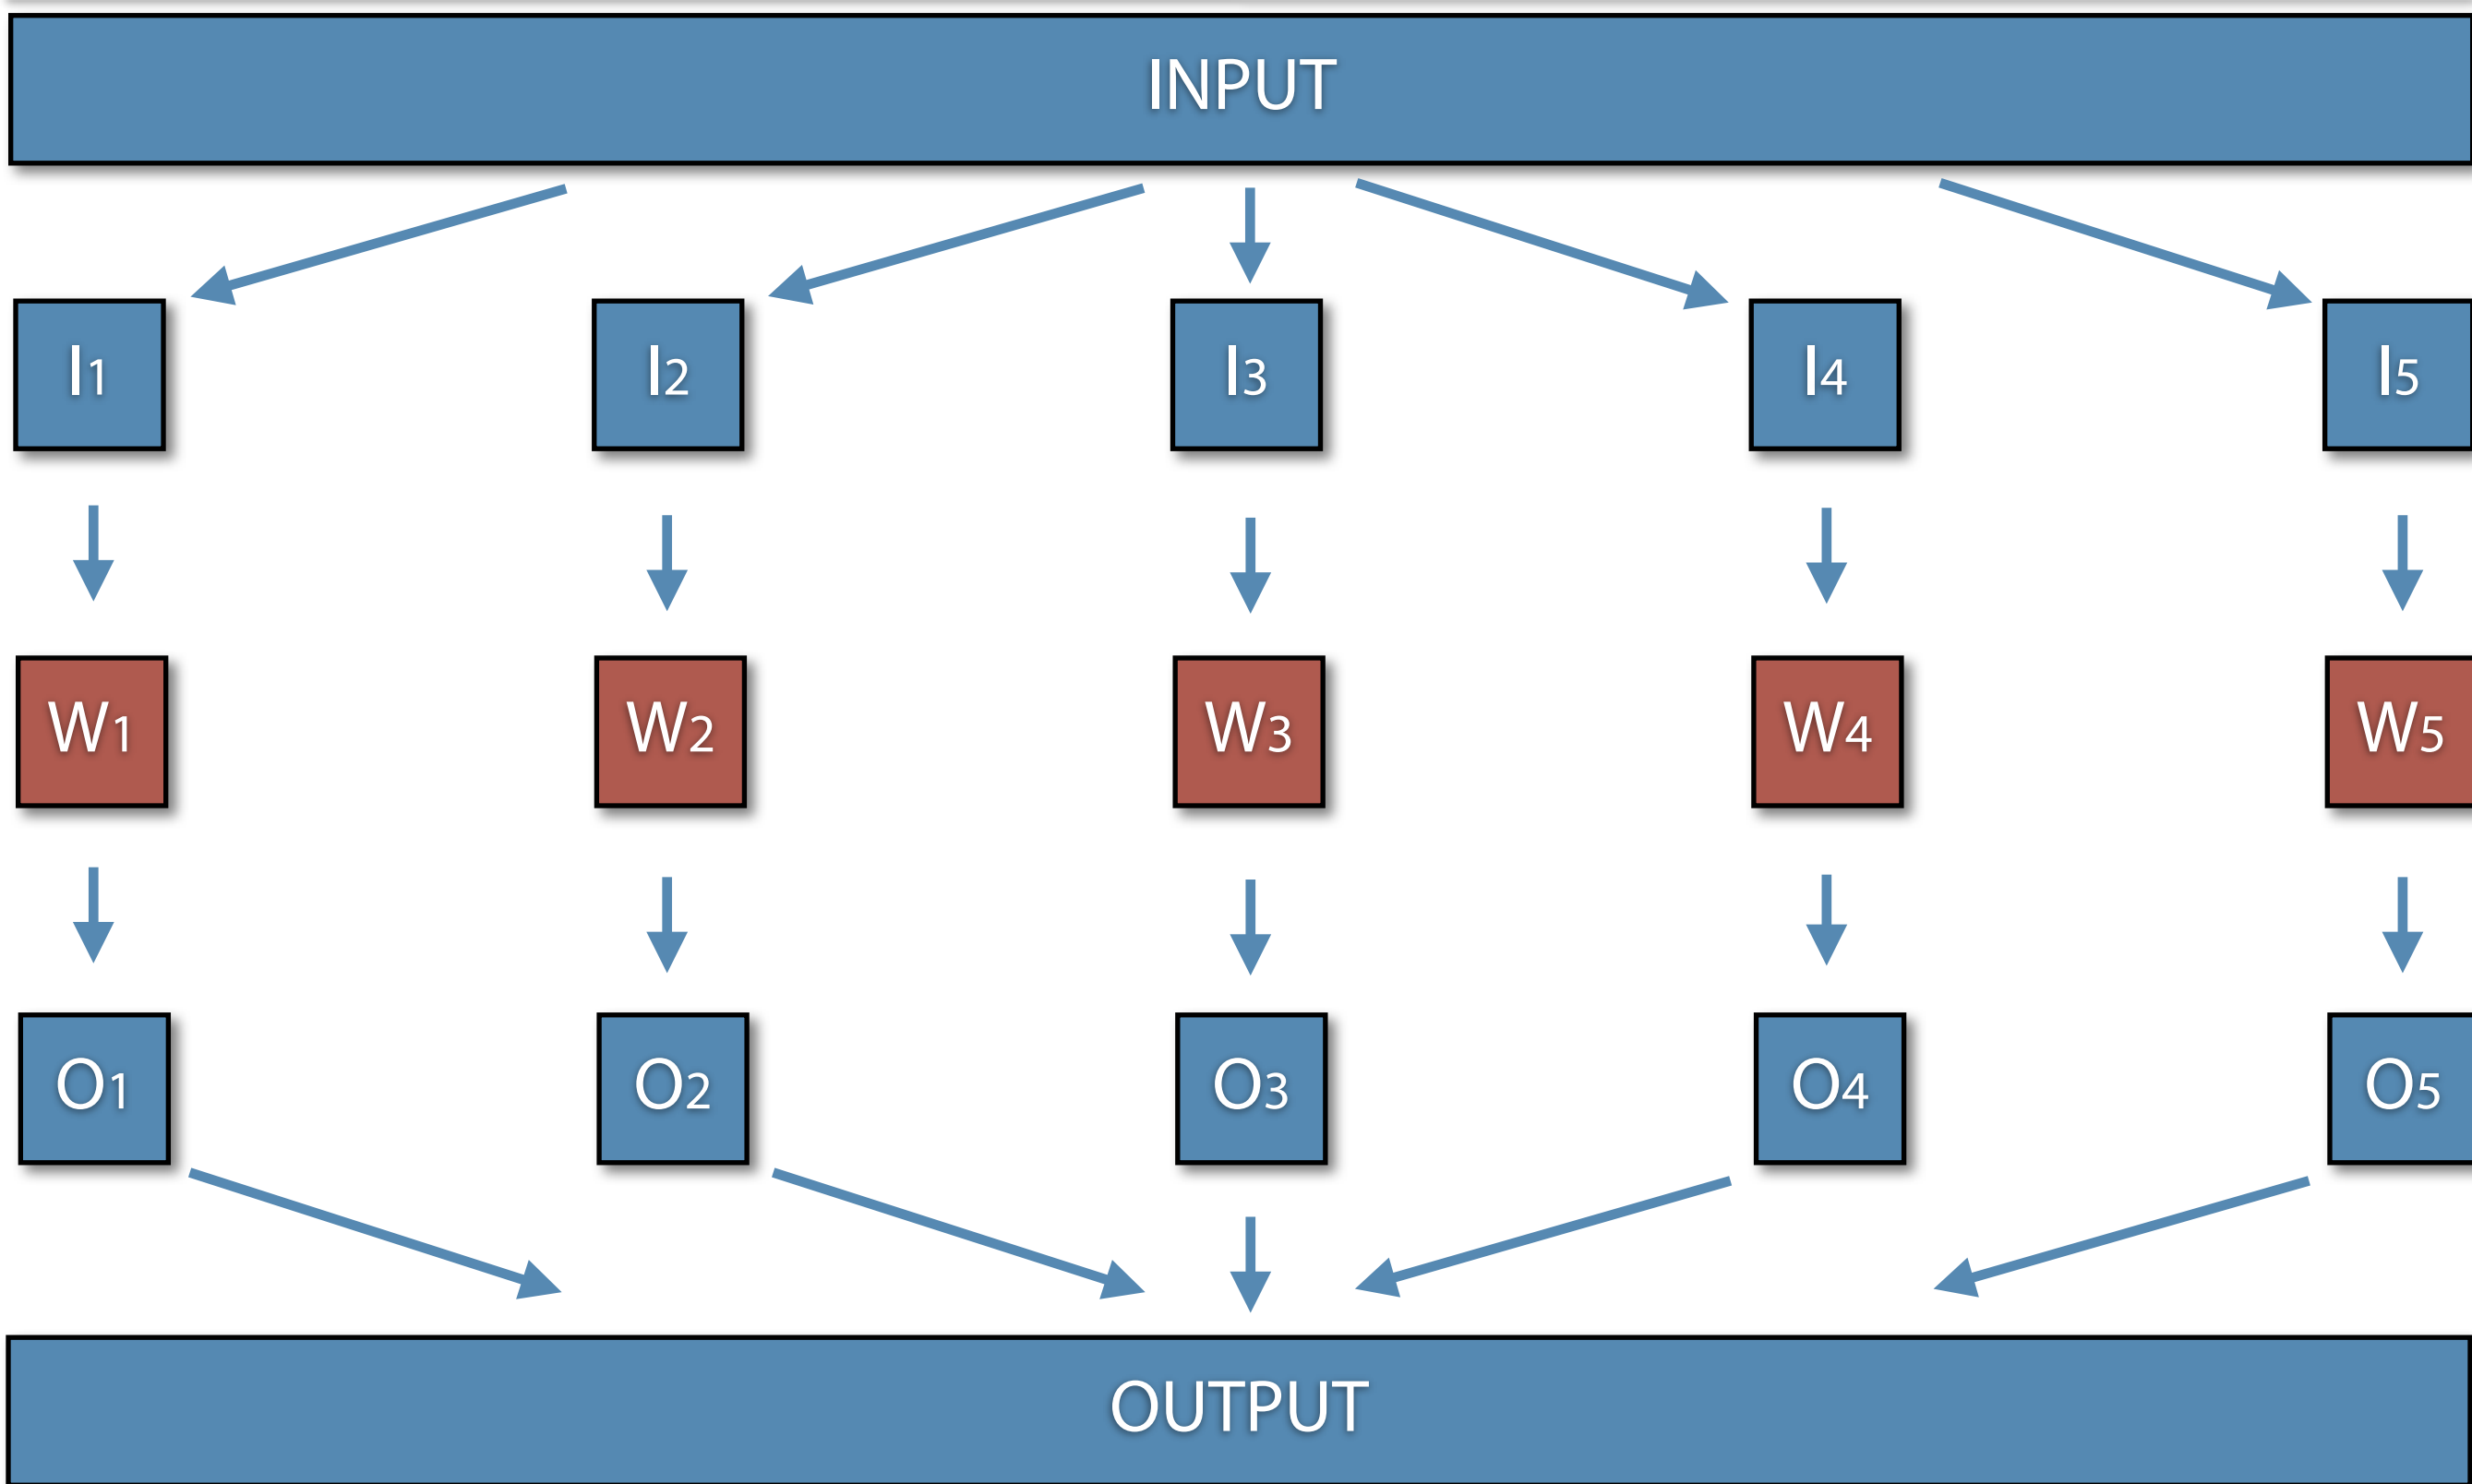


PROCESS



OUTPUT

Divide and Conquer



Questions

- How do we split the input?
- How do we distribute the input splits?
- How do we collect the output splits?
- How do we aggregate the output?
- How do we coordinate the work?
- What if input splits $>$ num workers?
- What if workers need to share input/output splits?
- What if a worker dies?
- What if we have a new input?

<http://www.duiops.net/seresvivos>

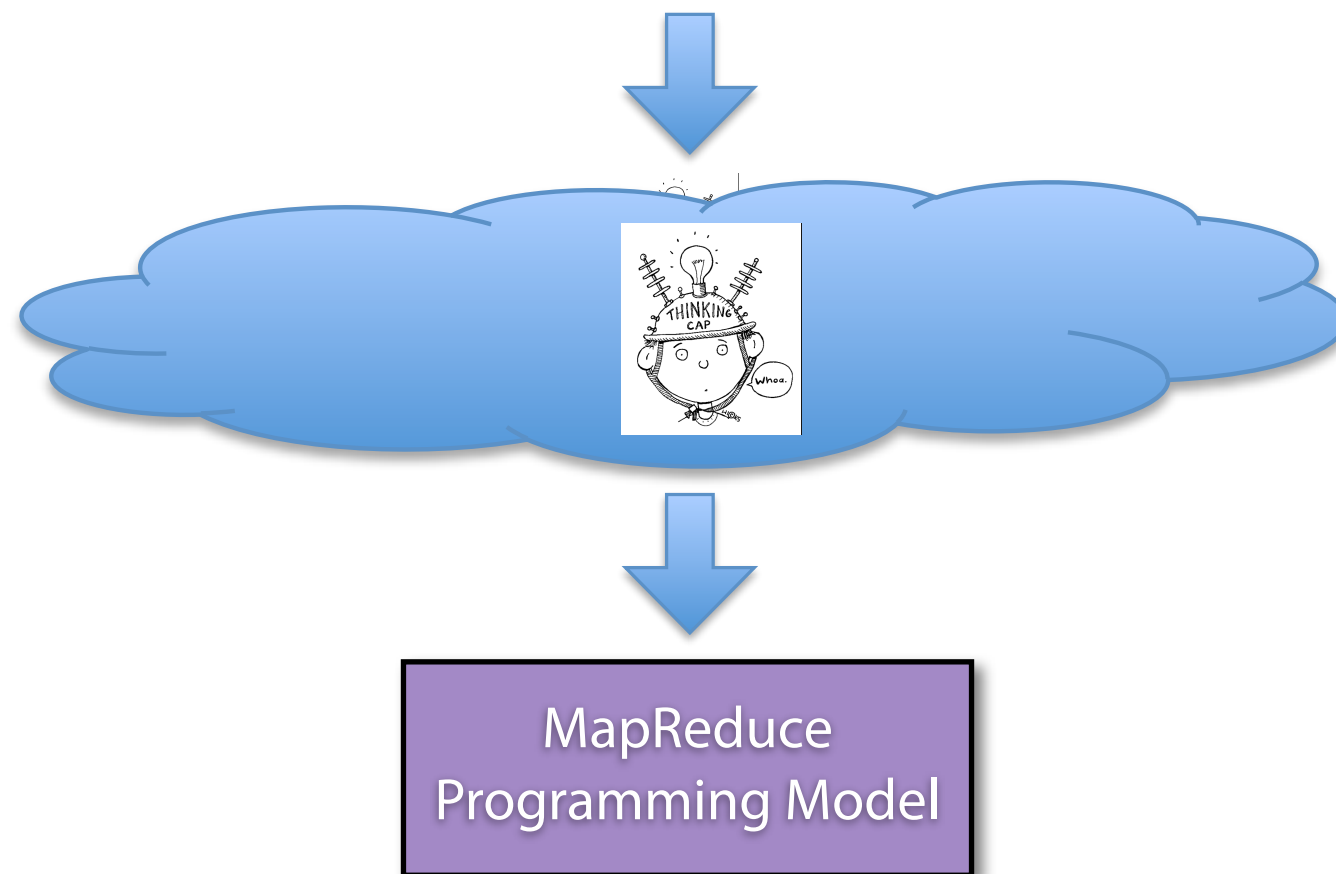


Design Ideas

- **Scale “out”, not “up”**
 - Low end machines
- **Move processing to the data**
 - Network bandwidth bottleneck
- **Process data sequentially, avoid random access**
 - Huge data files
 - Write once, read many
- **Seamless scalability**
 - Strive for the unobtainable
- **Right level of abstraction**
 - Hide implementation details from applications development

Typical Large-Data Problem


- Iterate over a large number of records
- Extract something of interest from each
- Shuffle and sort intermediate results
- Aggregate intermediate results
- Generate final output




From functional programming...

From functional programming...


Input List



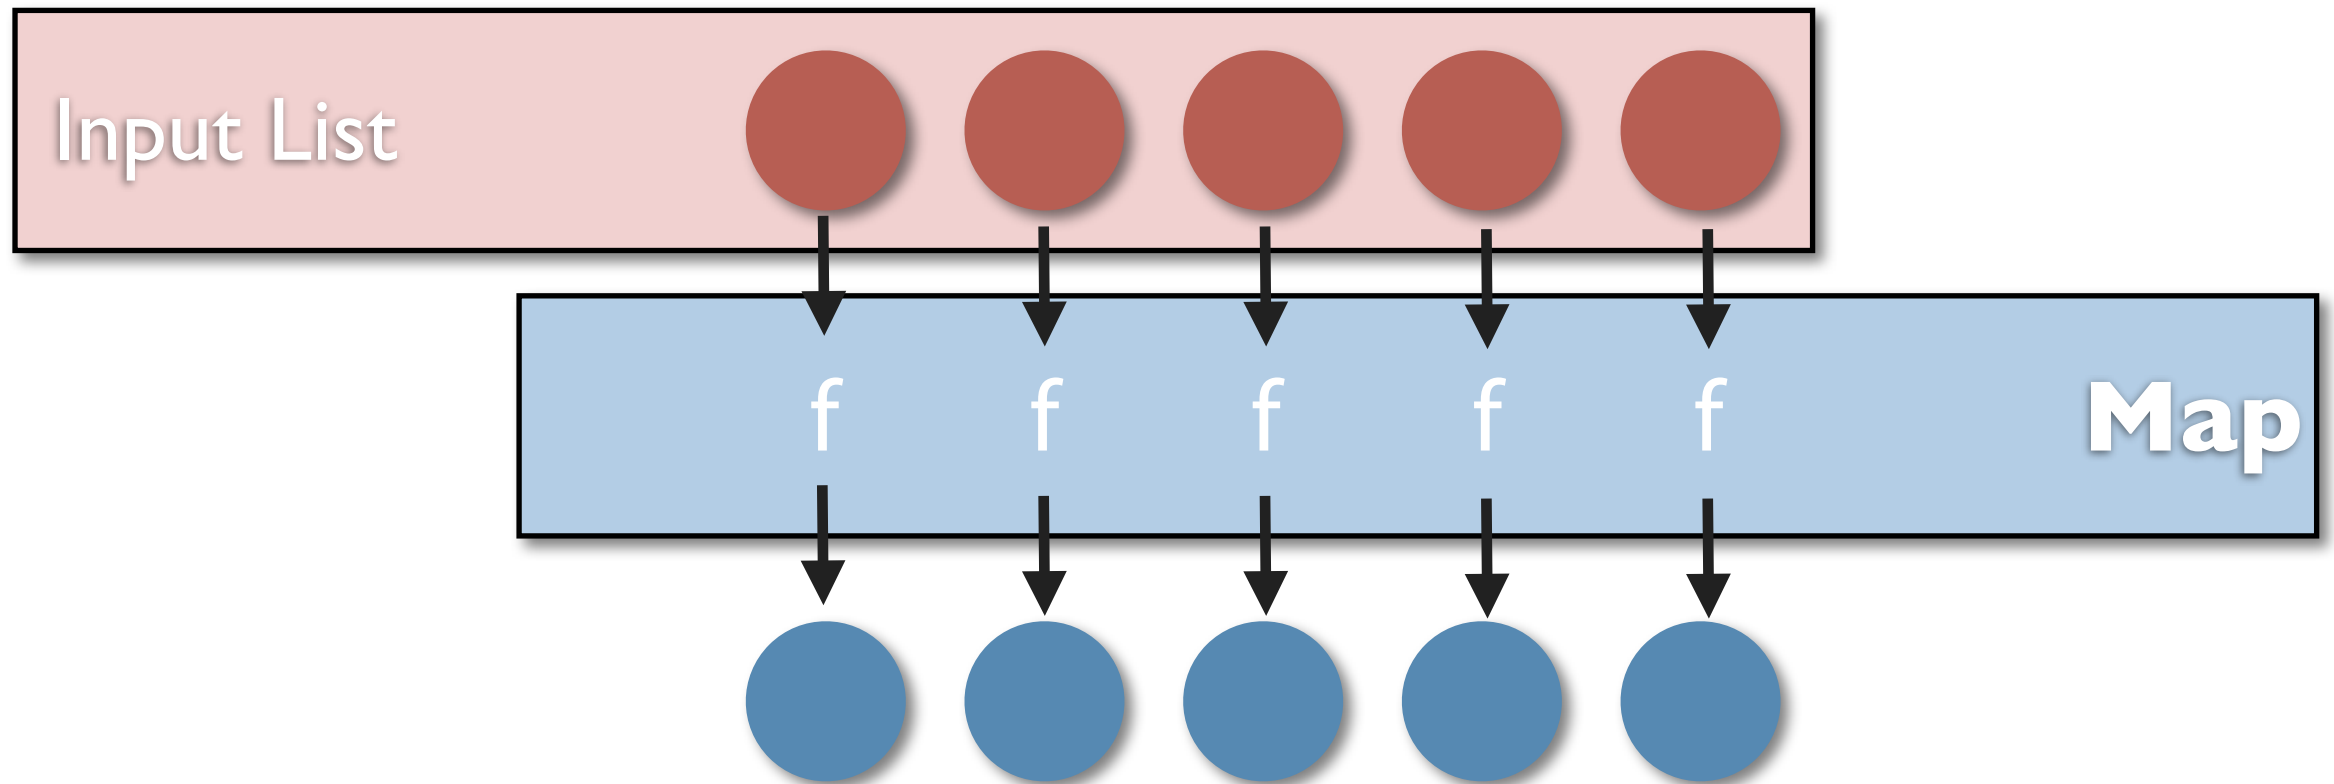
From functional programming...




From functional programming...



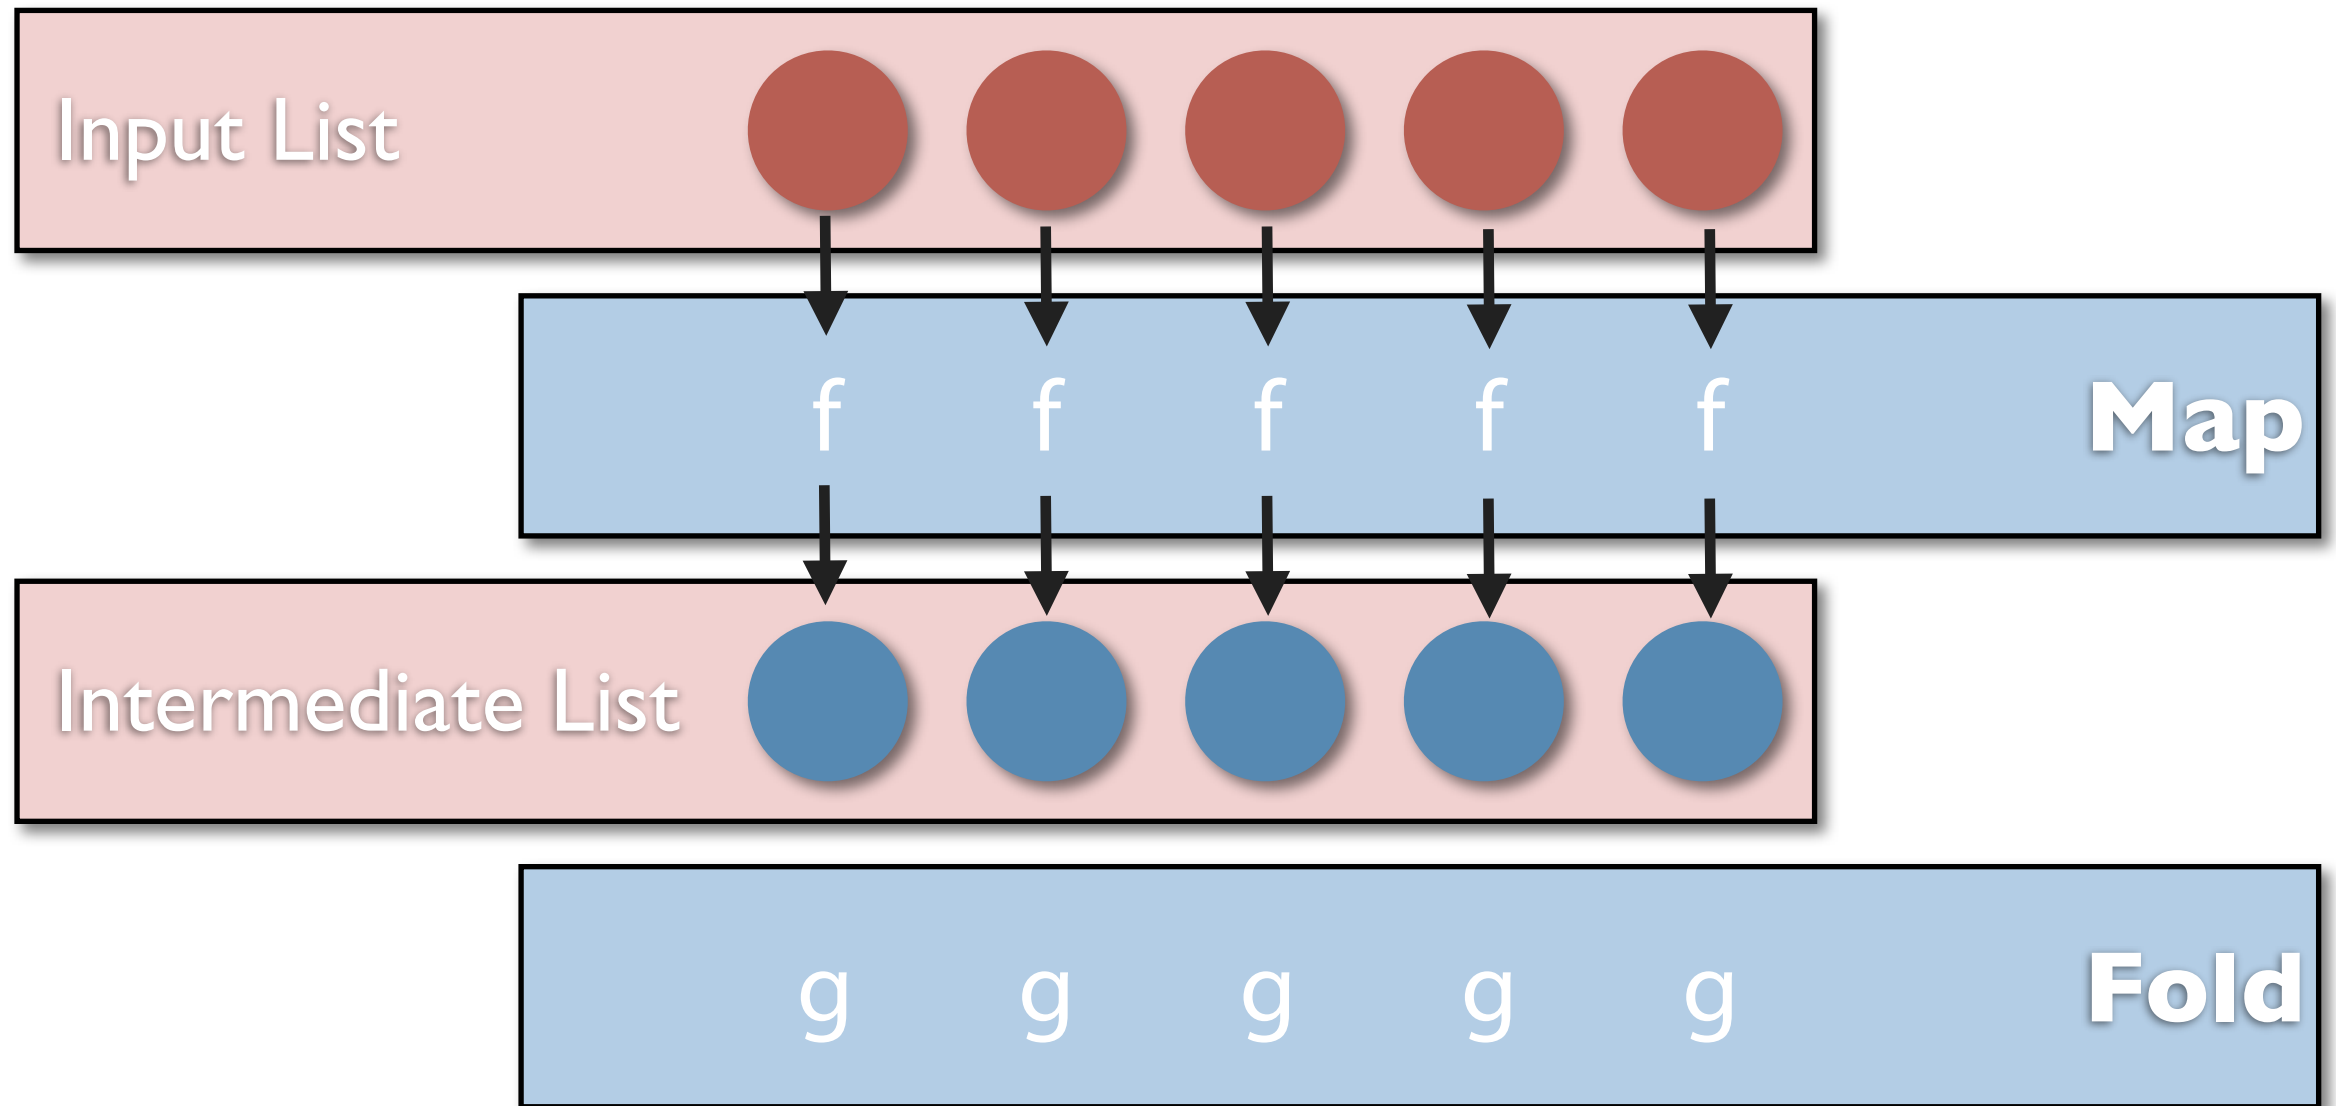
From functional programming...



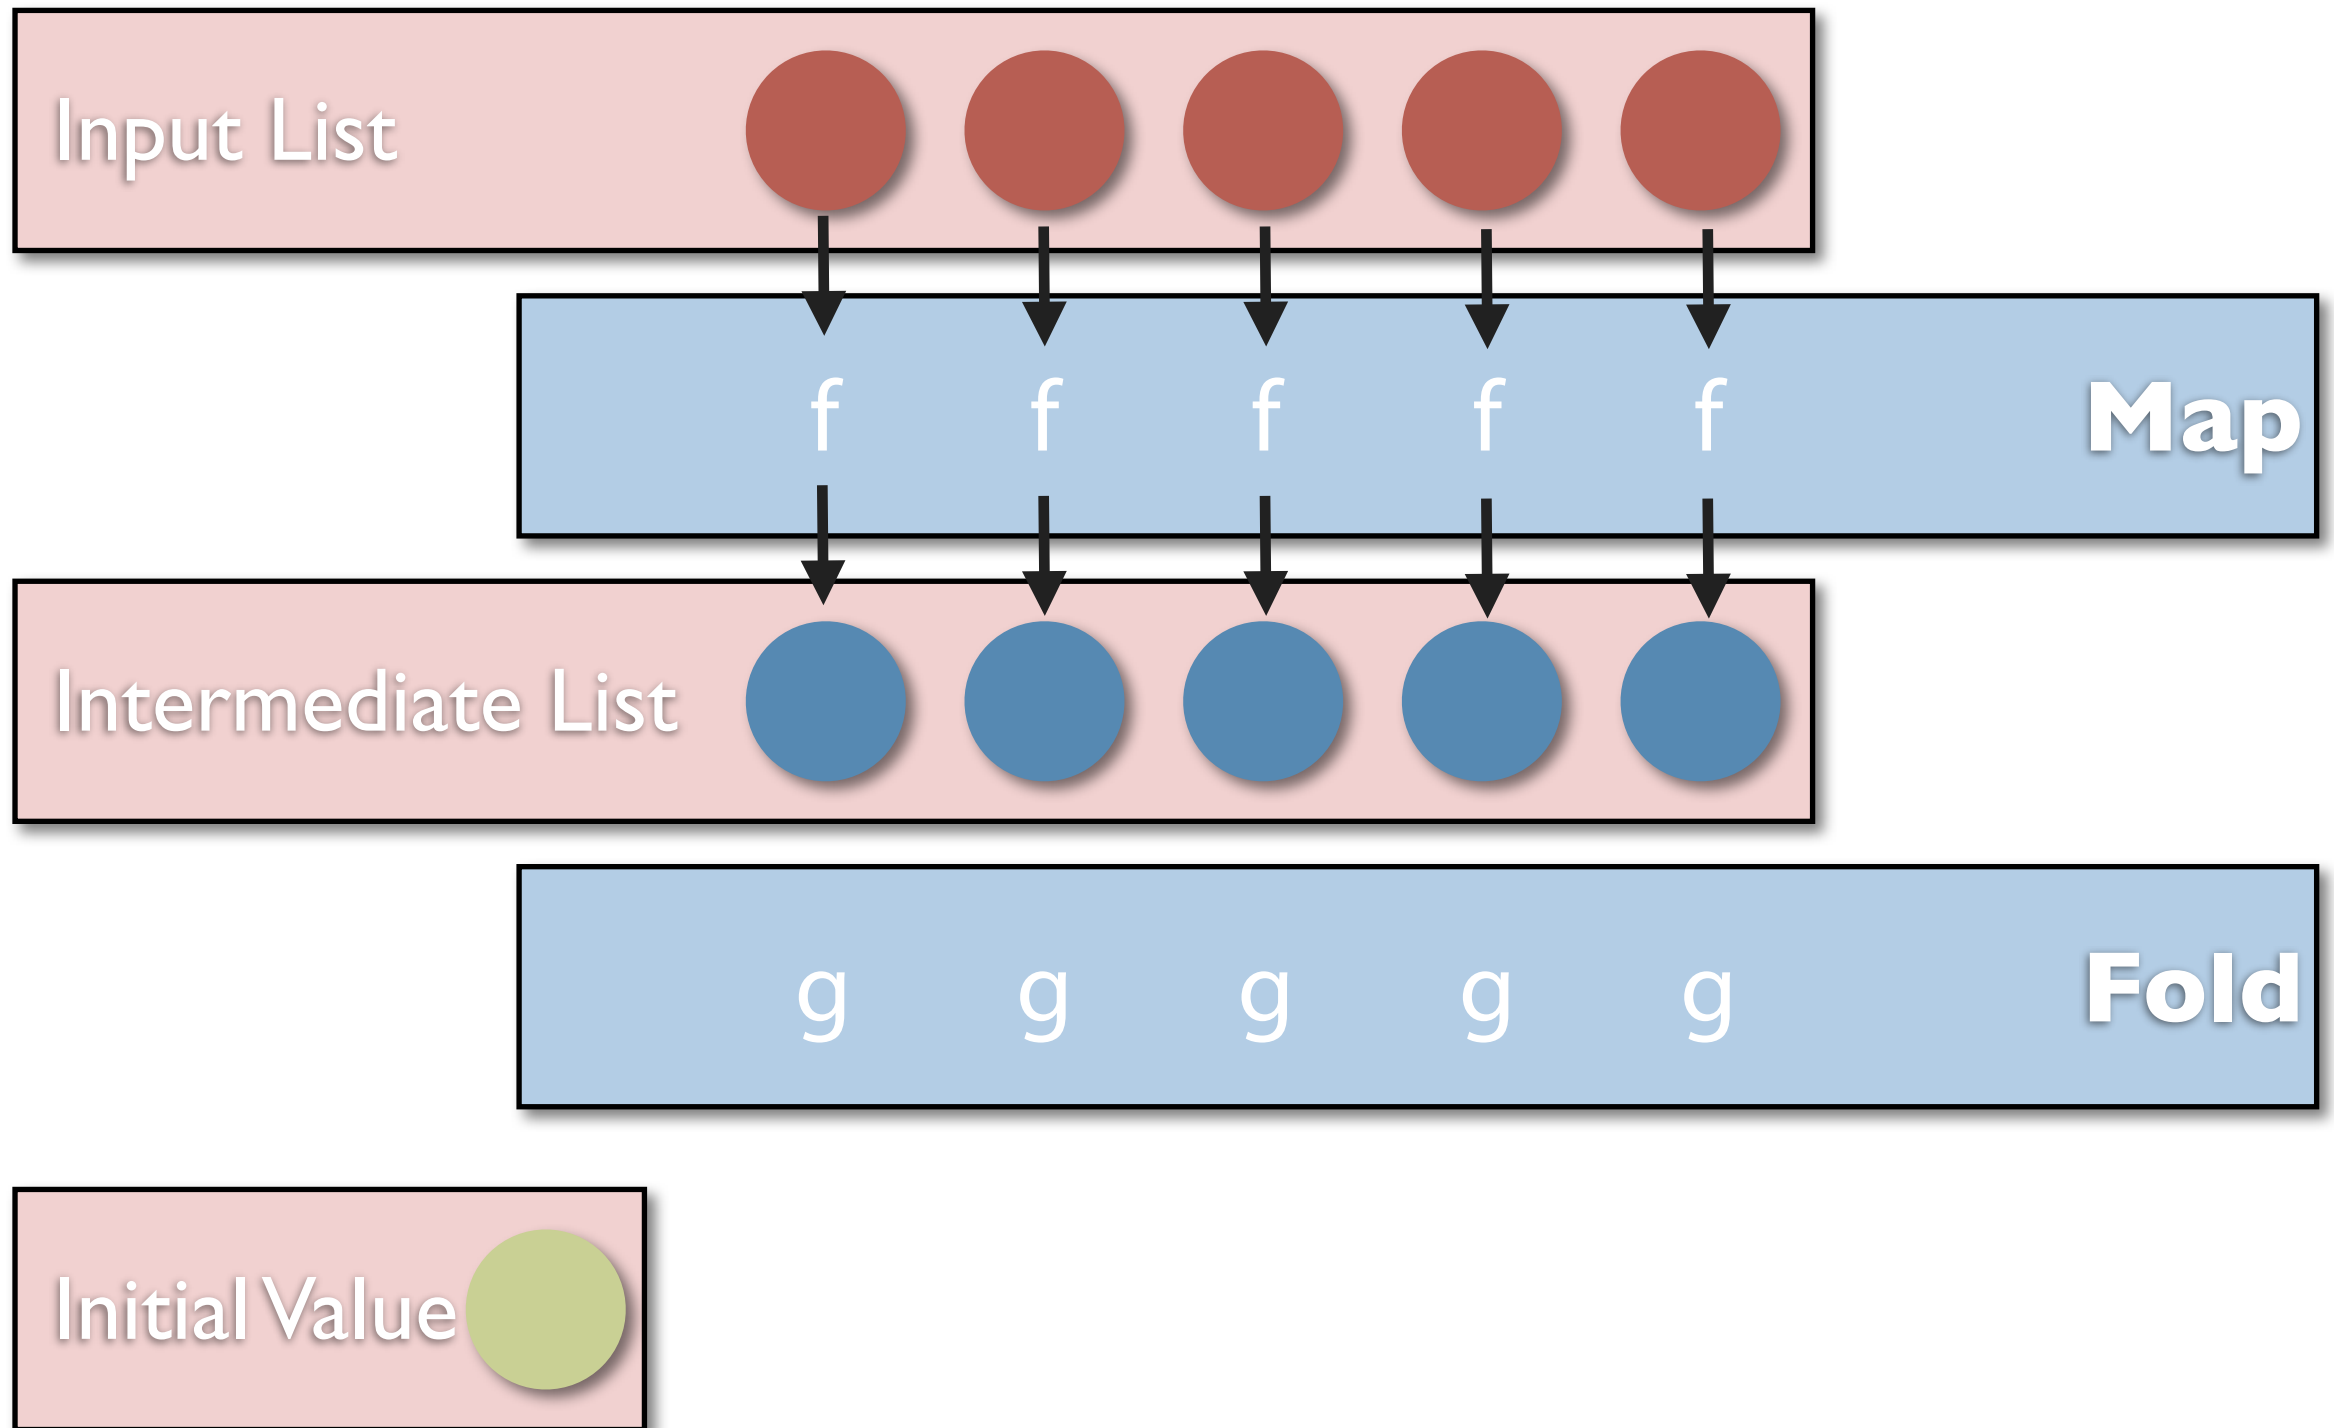
From functional programming...



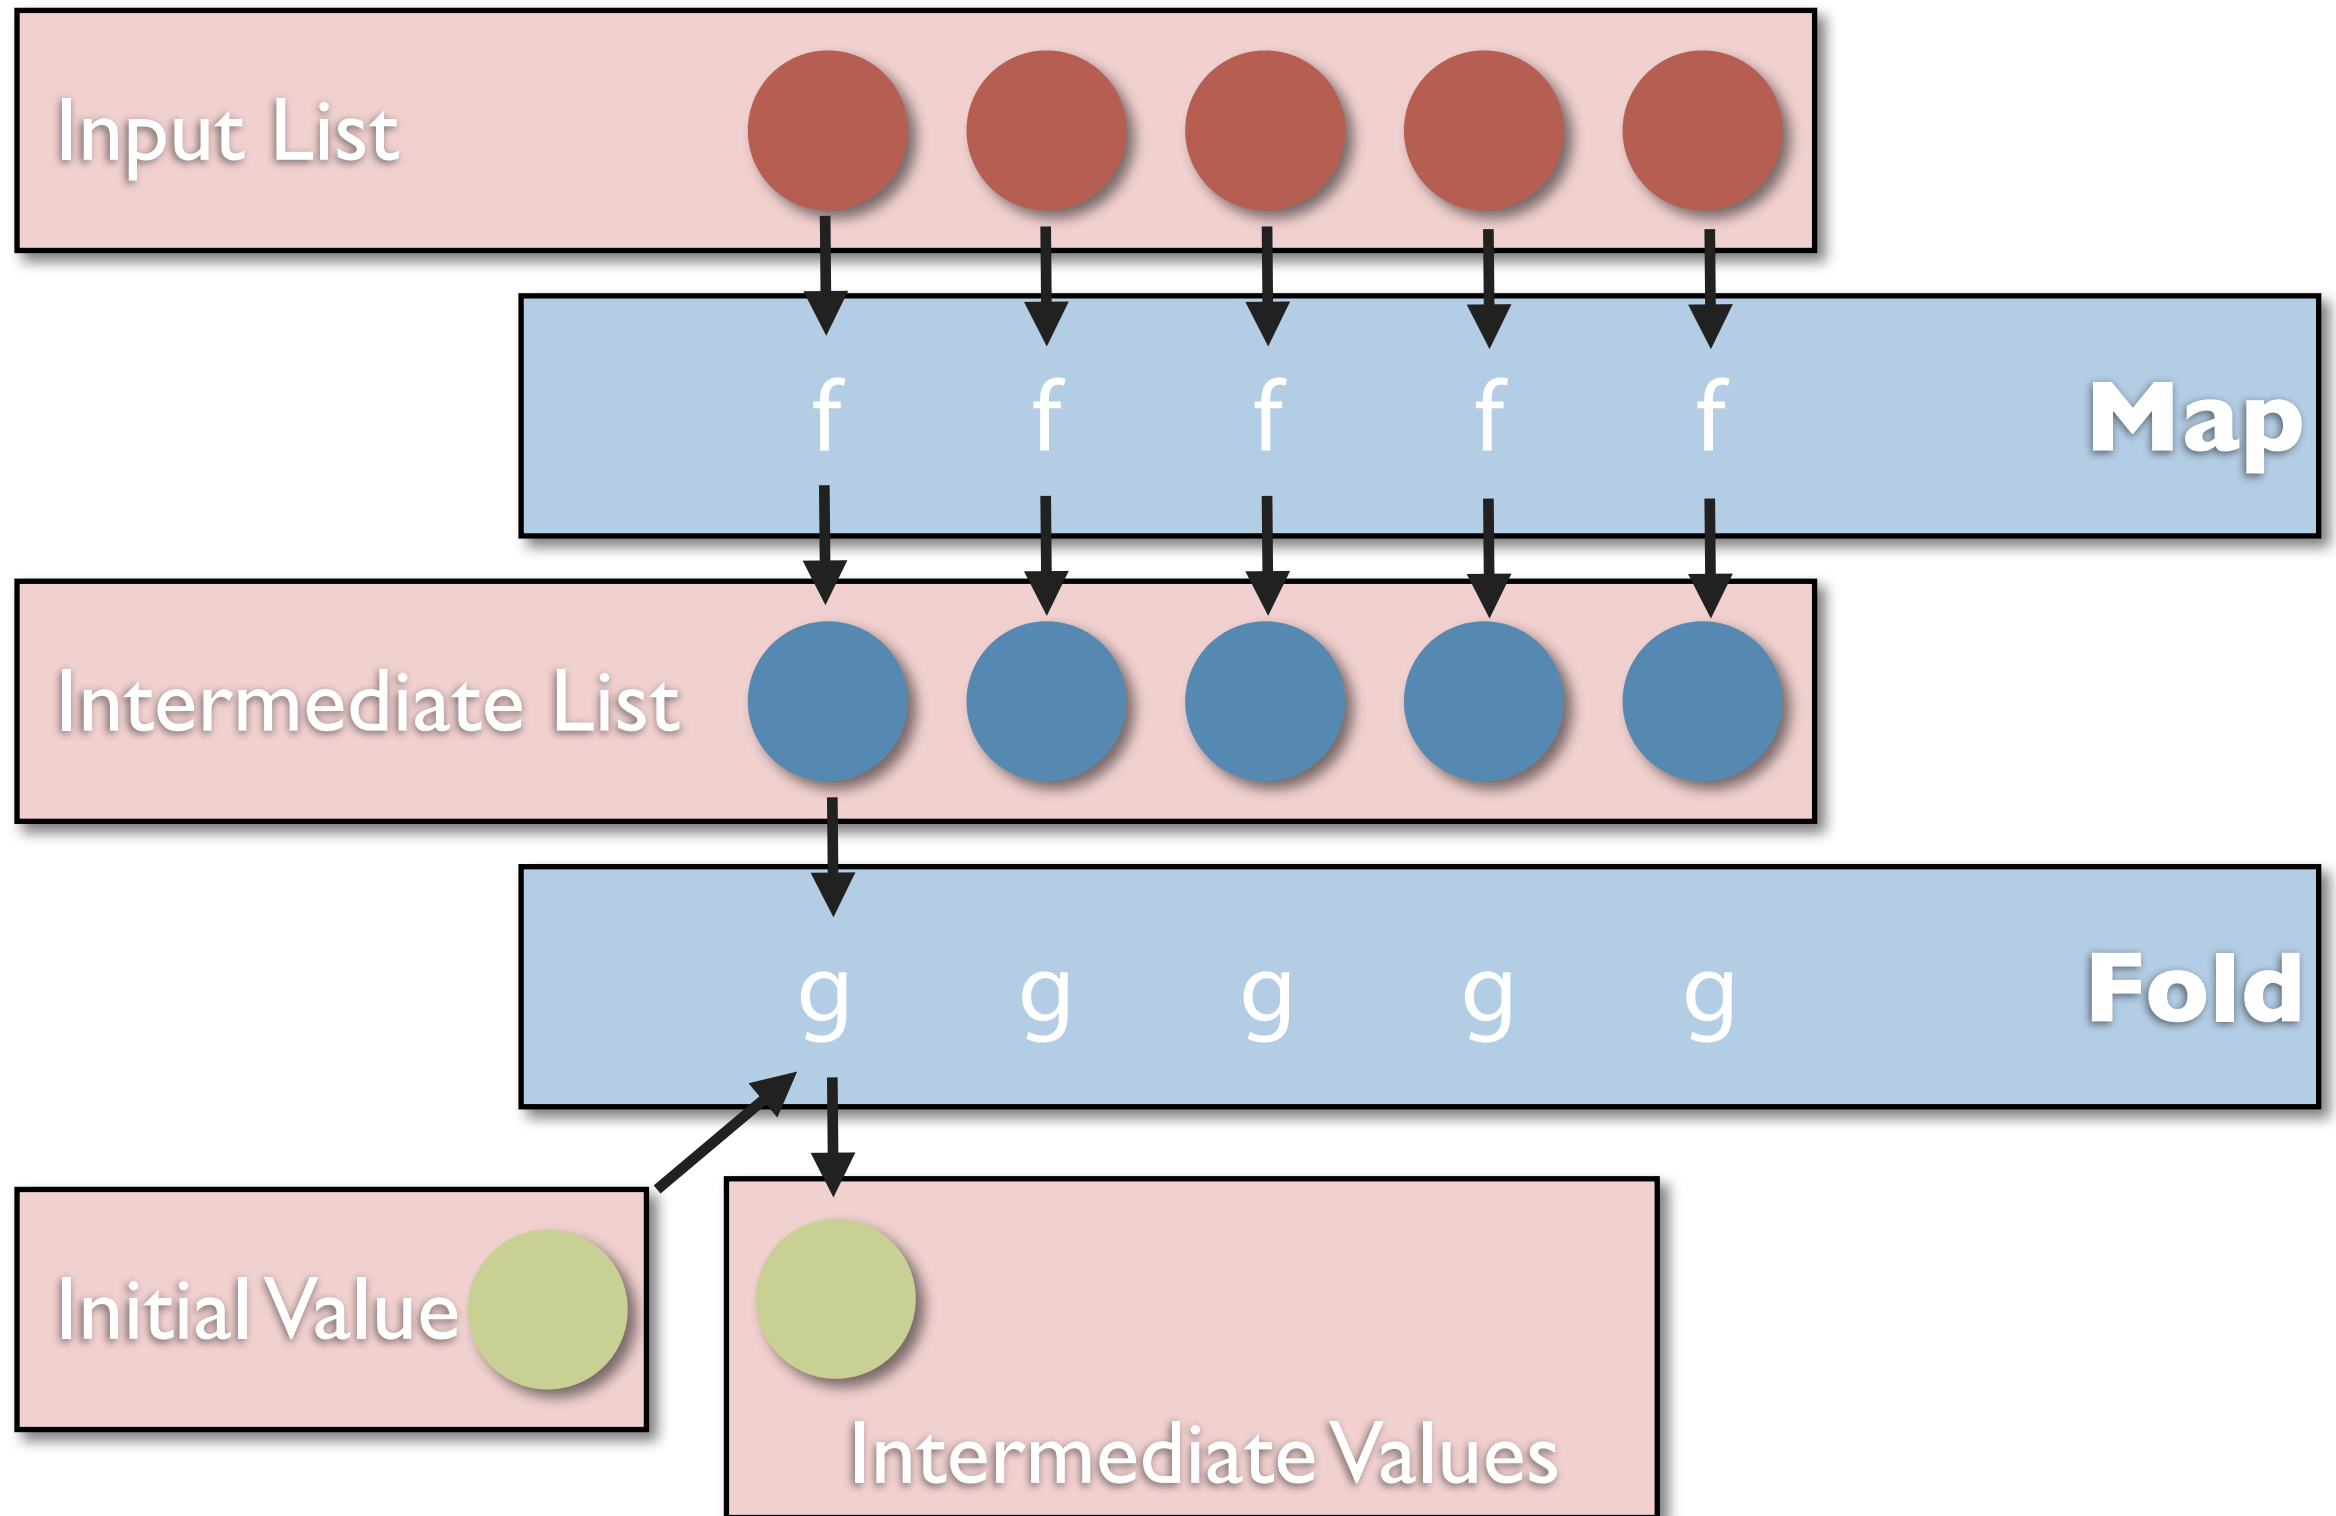
From functional programming...



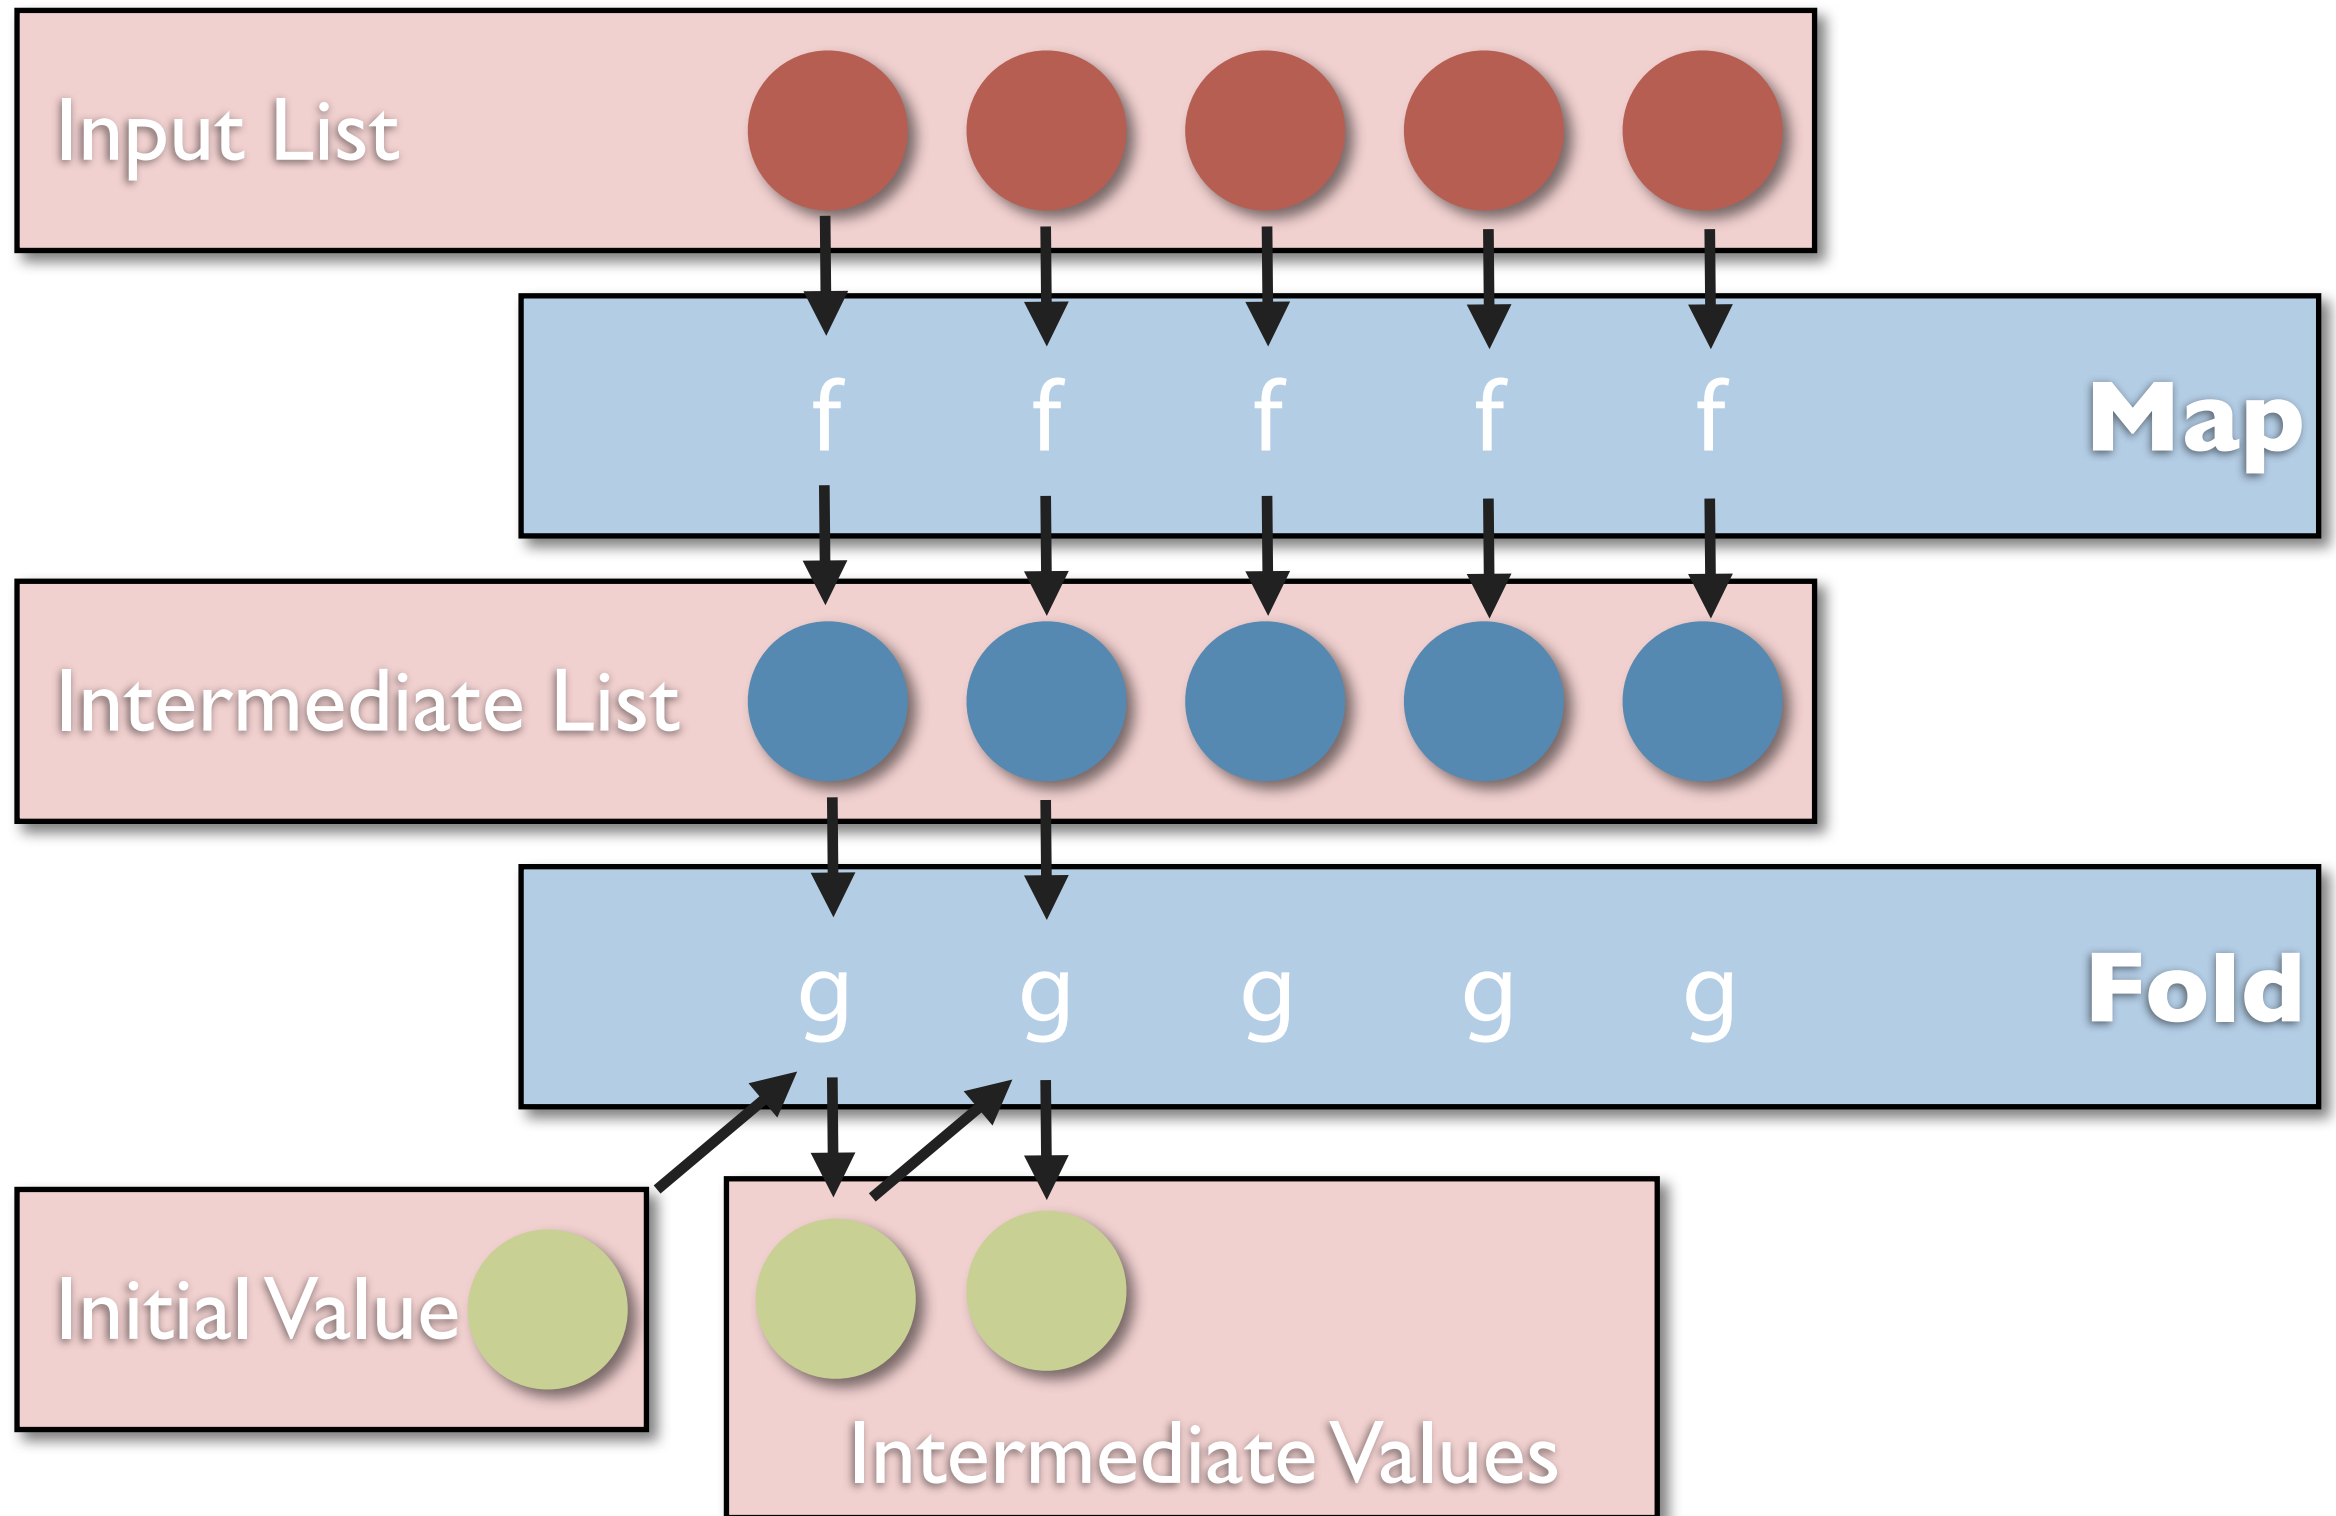
From functional programming...



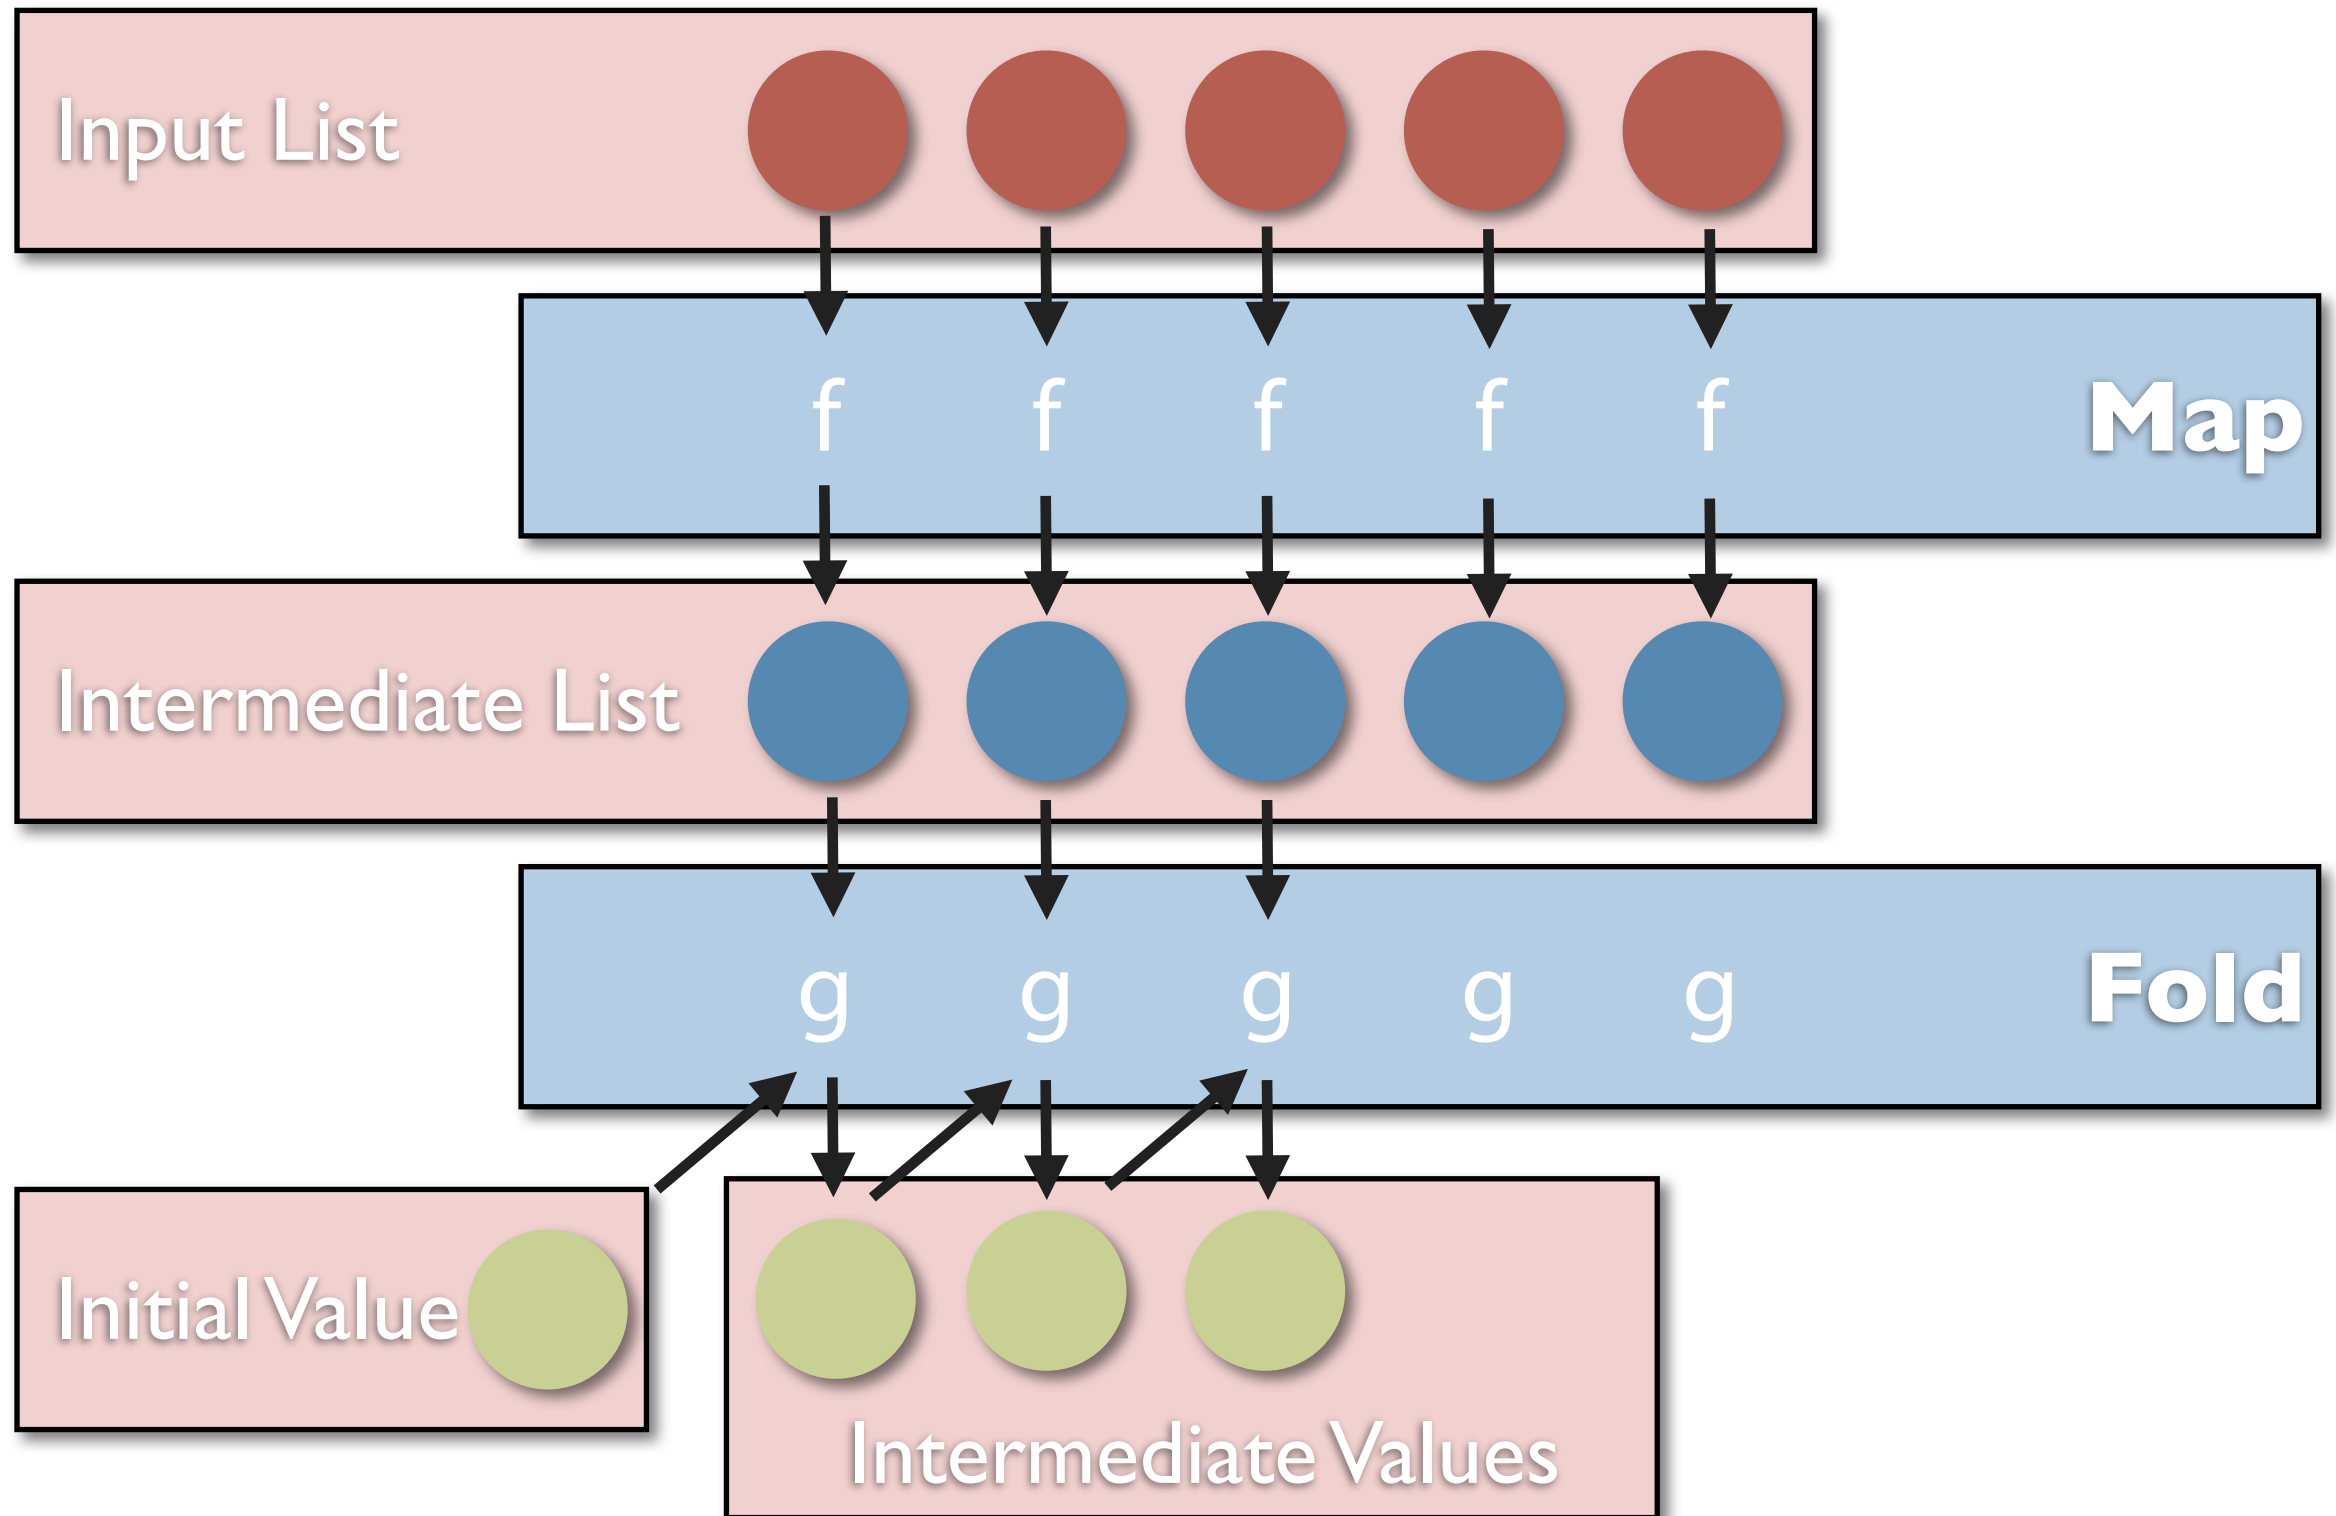
From functional programming...



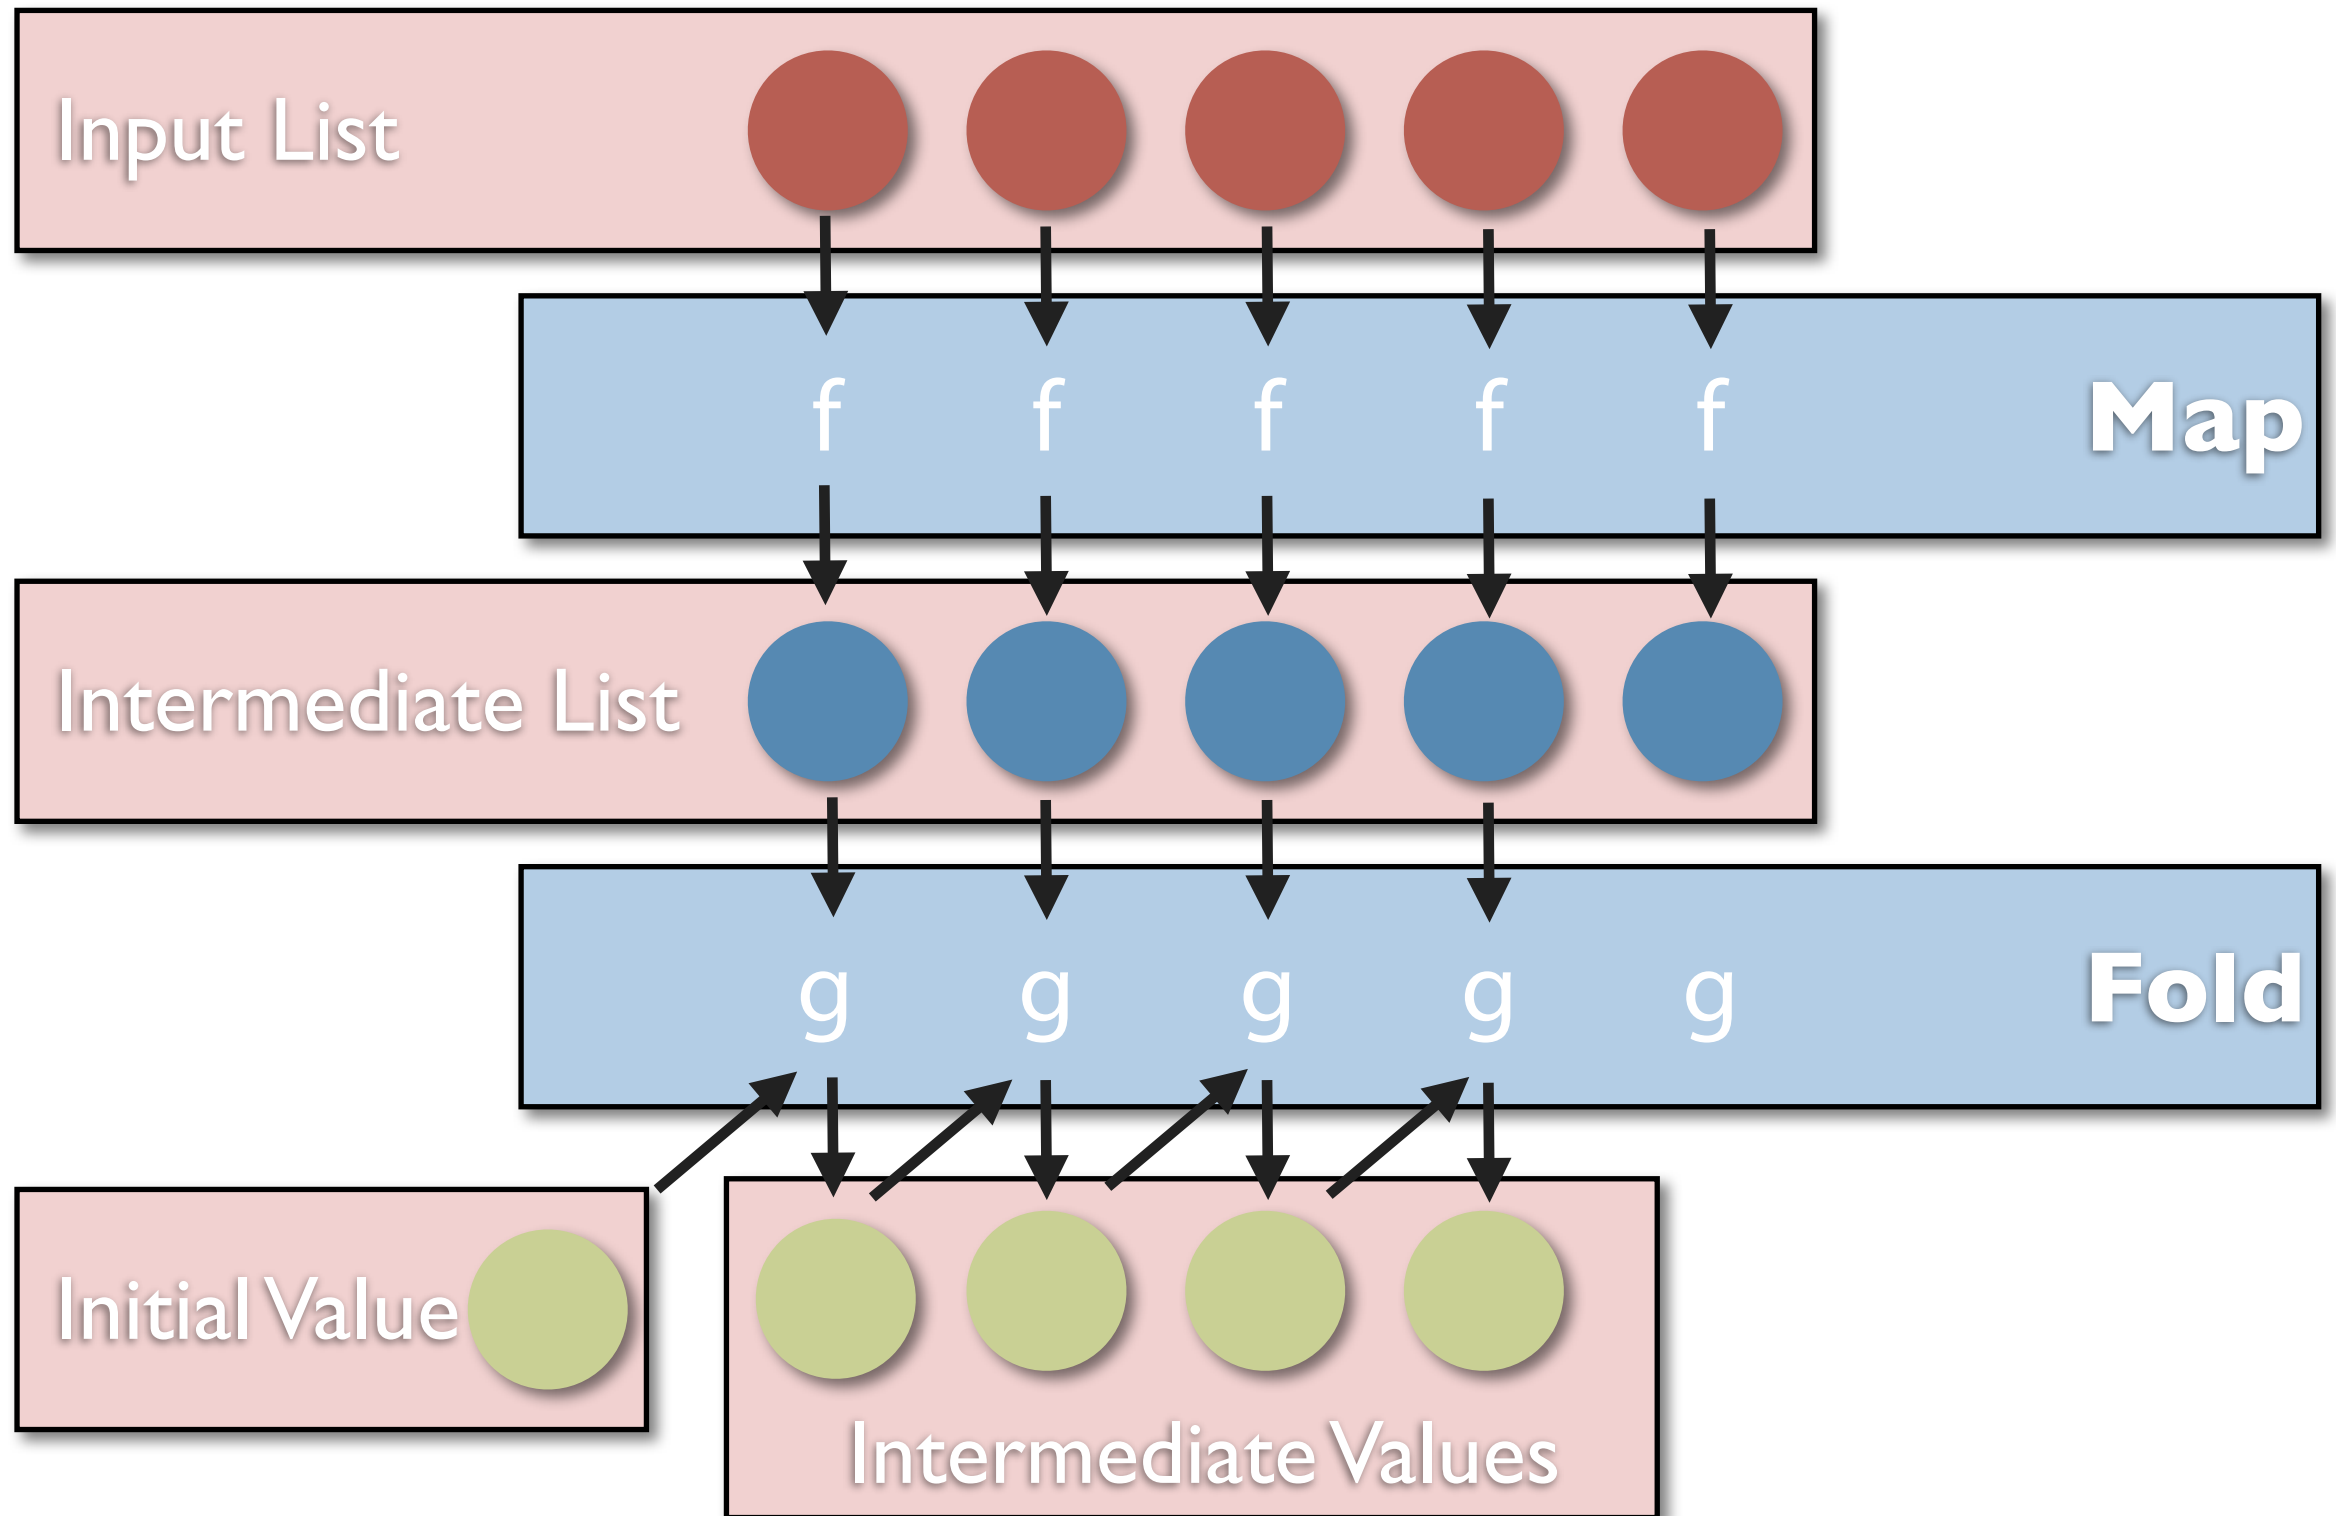
From functional programming...



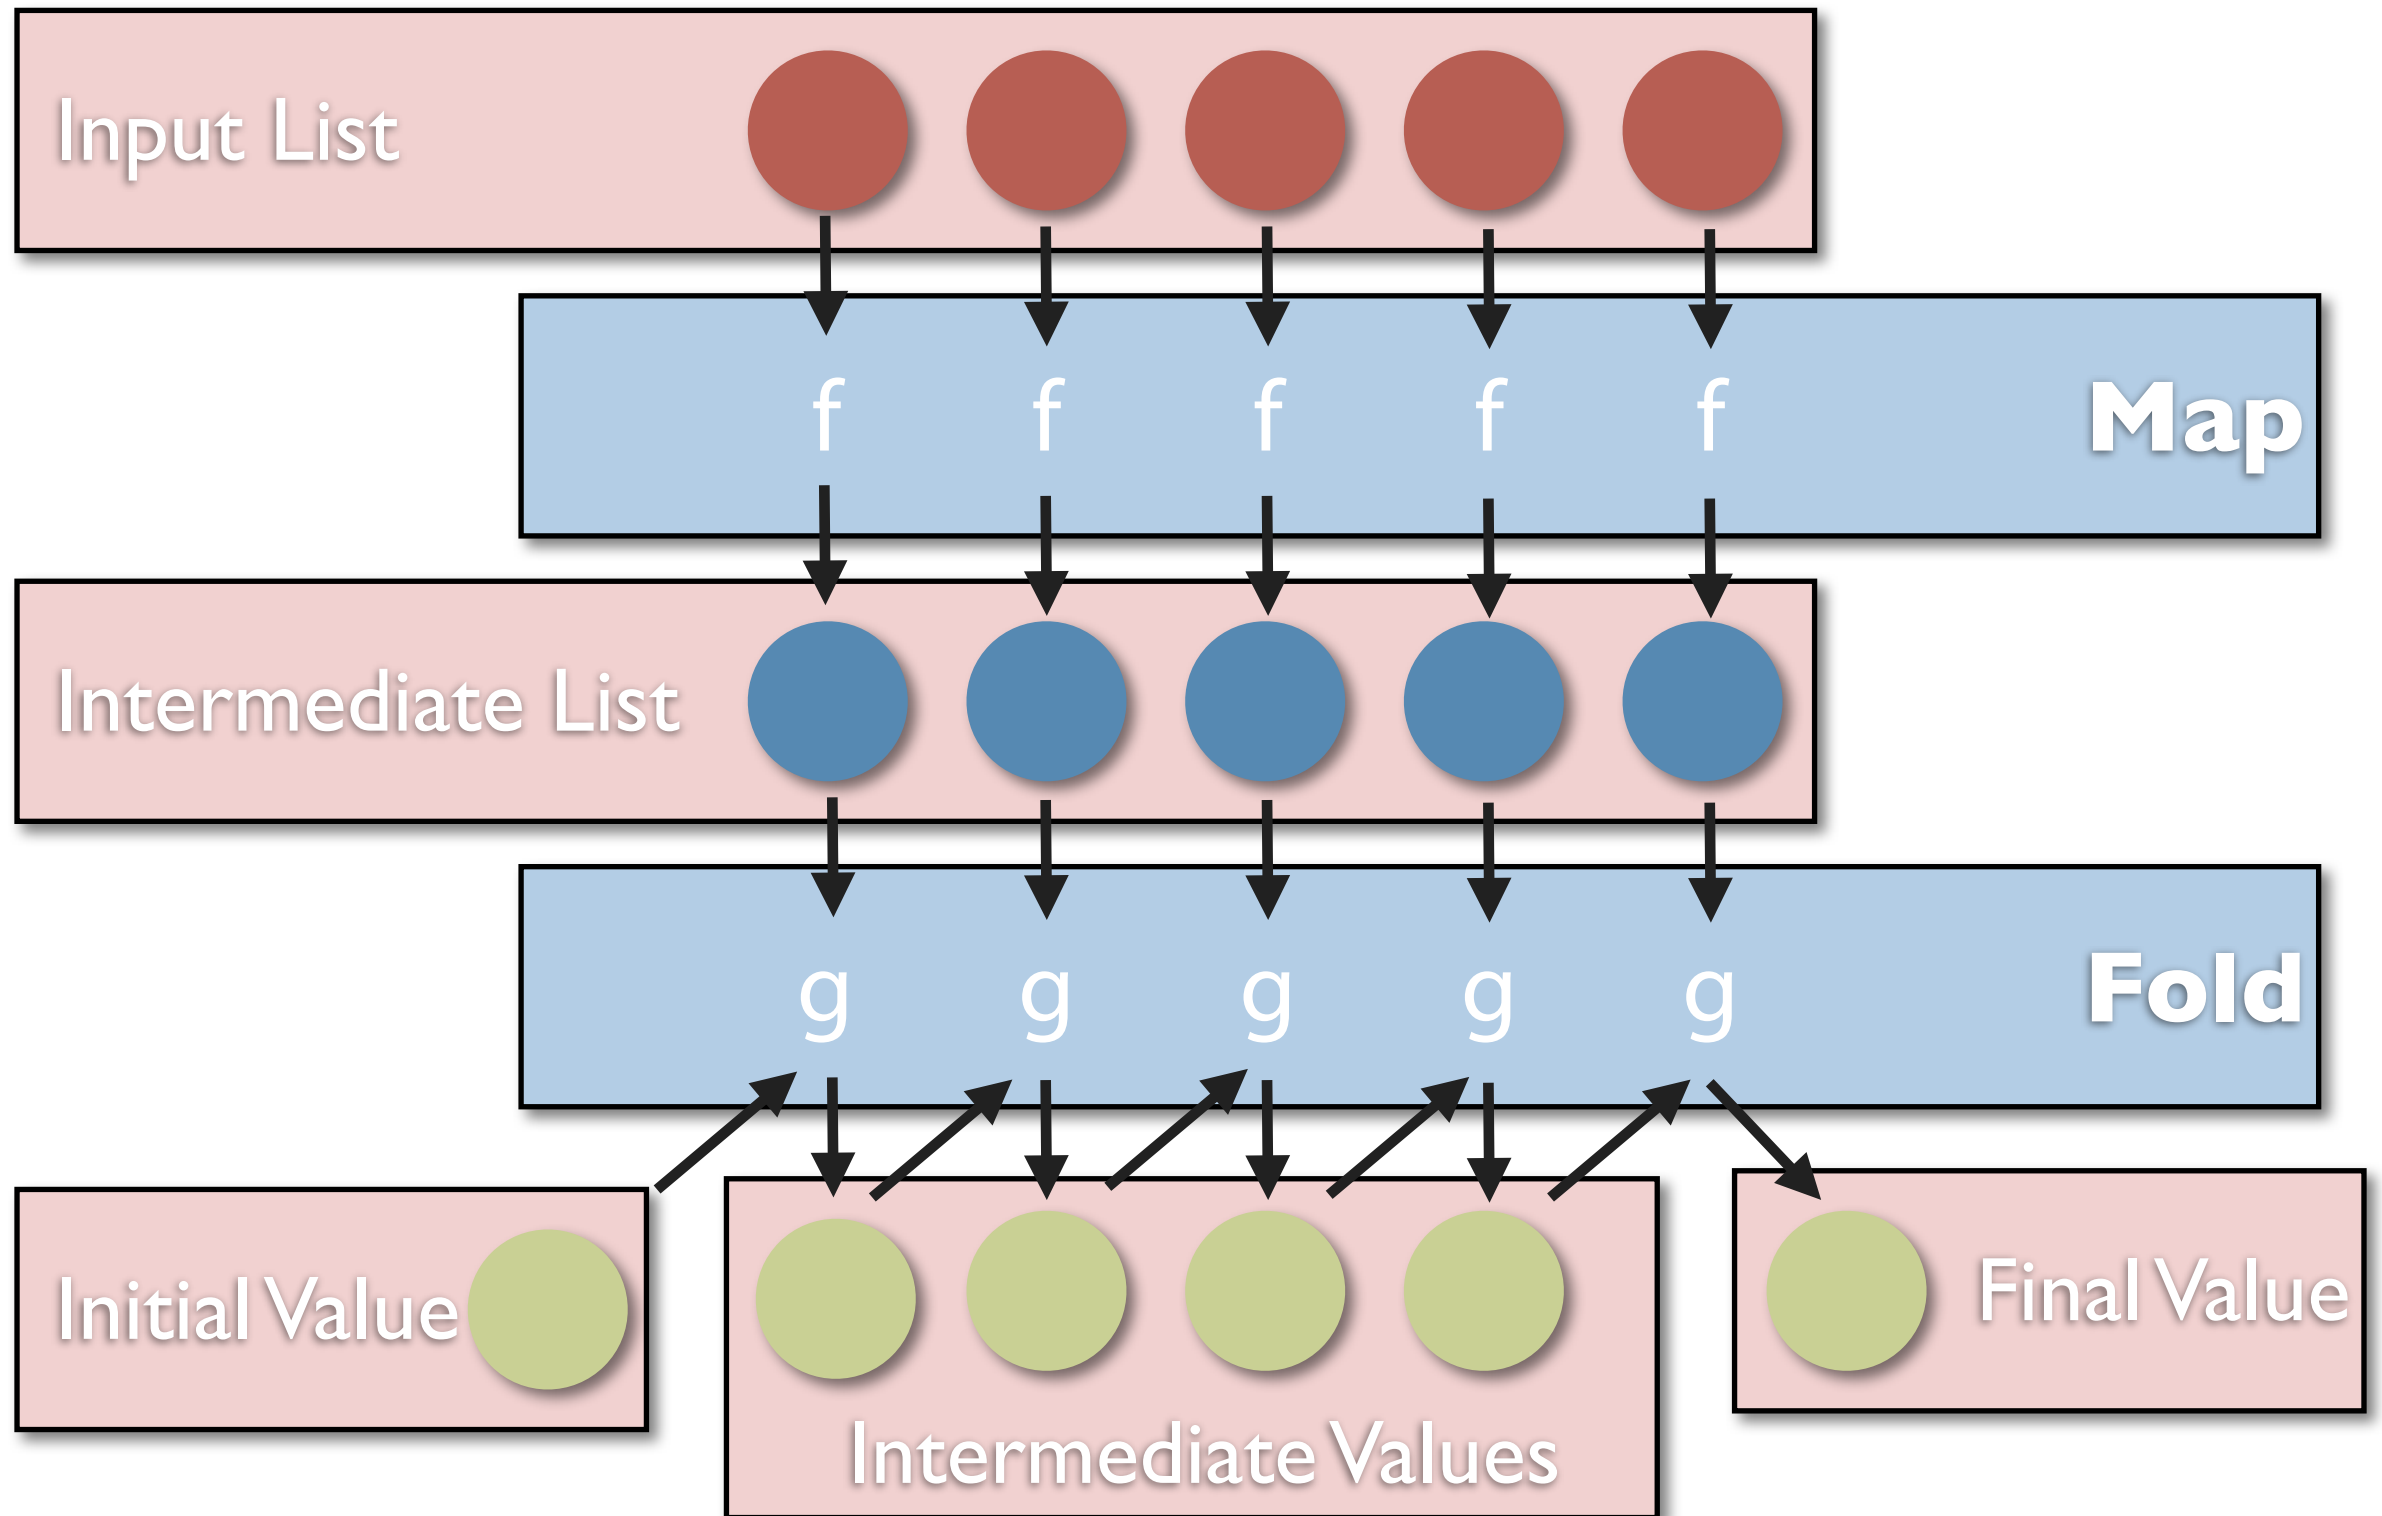
From functional programming...



From functional programming...



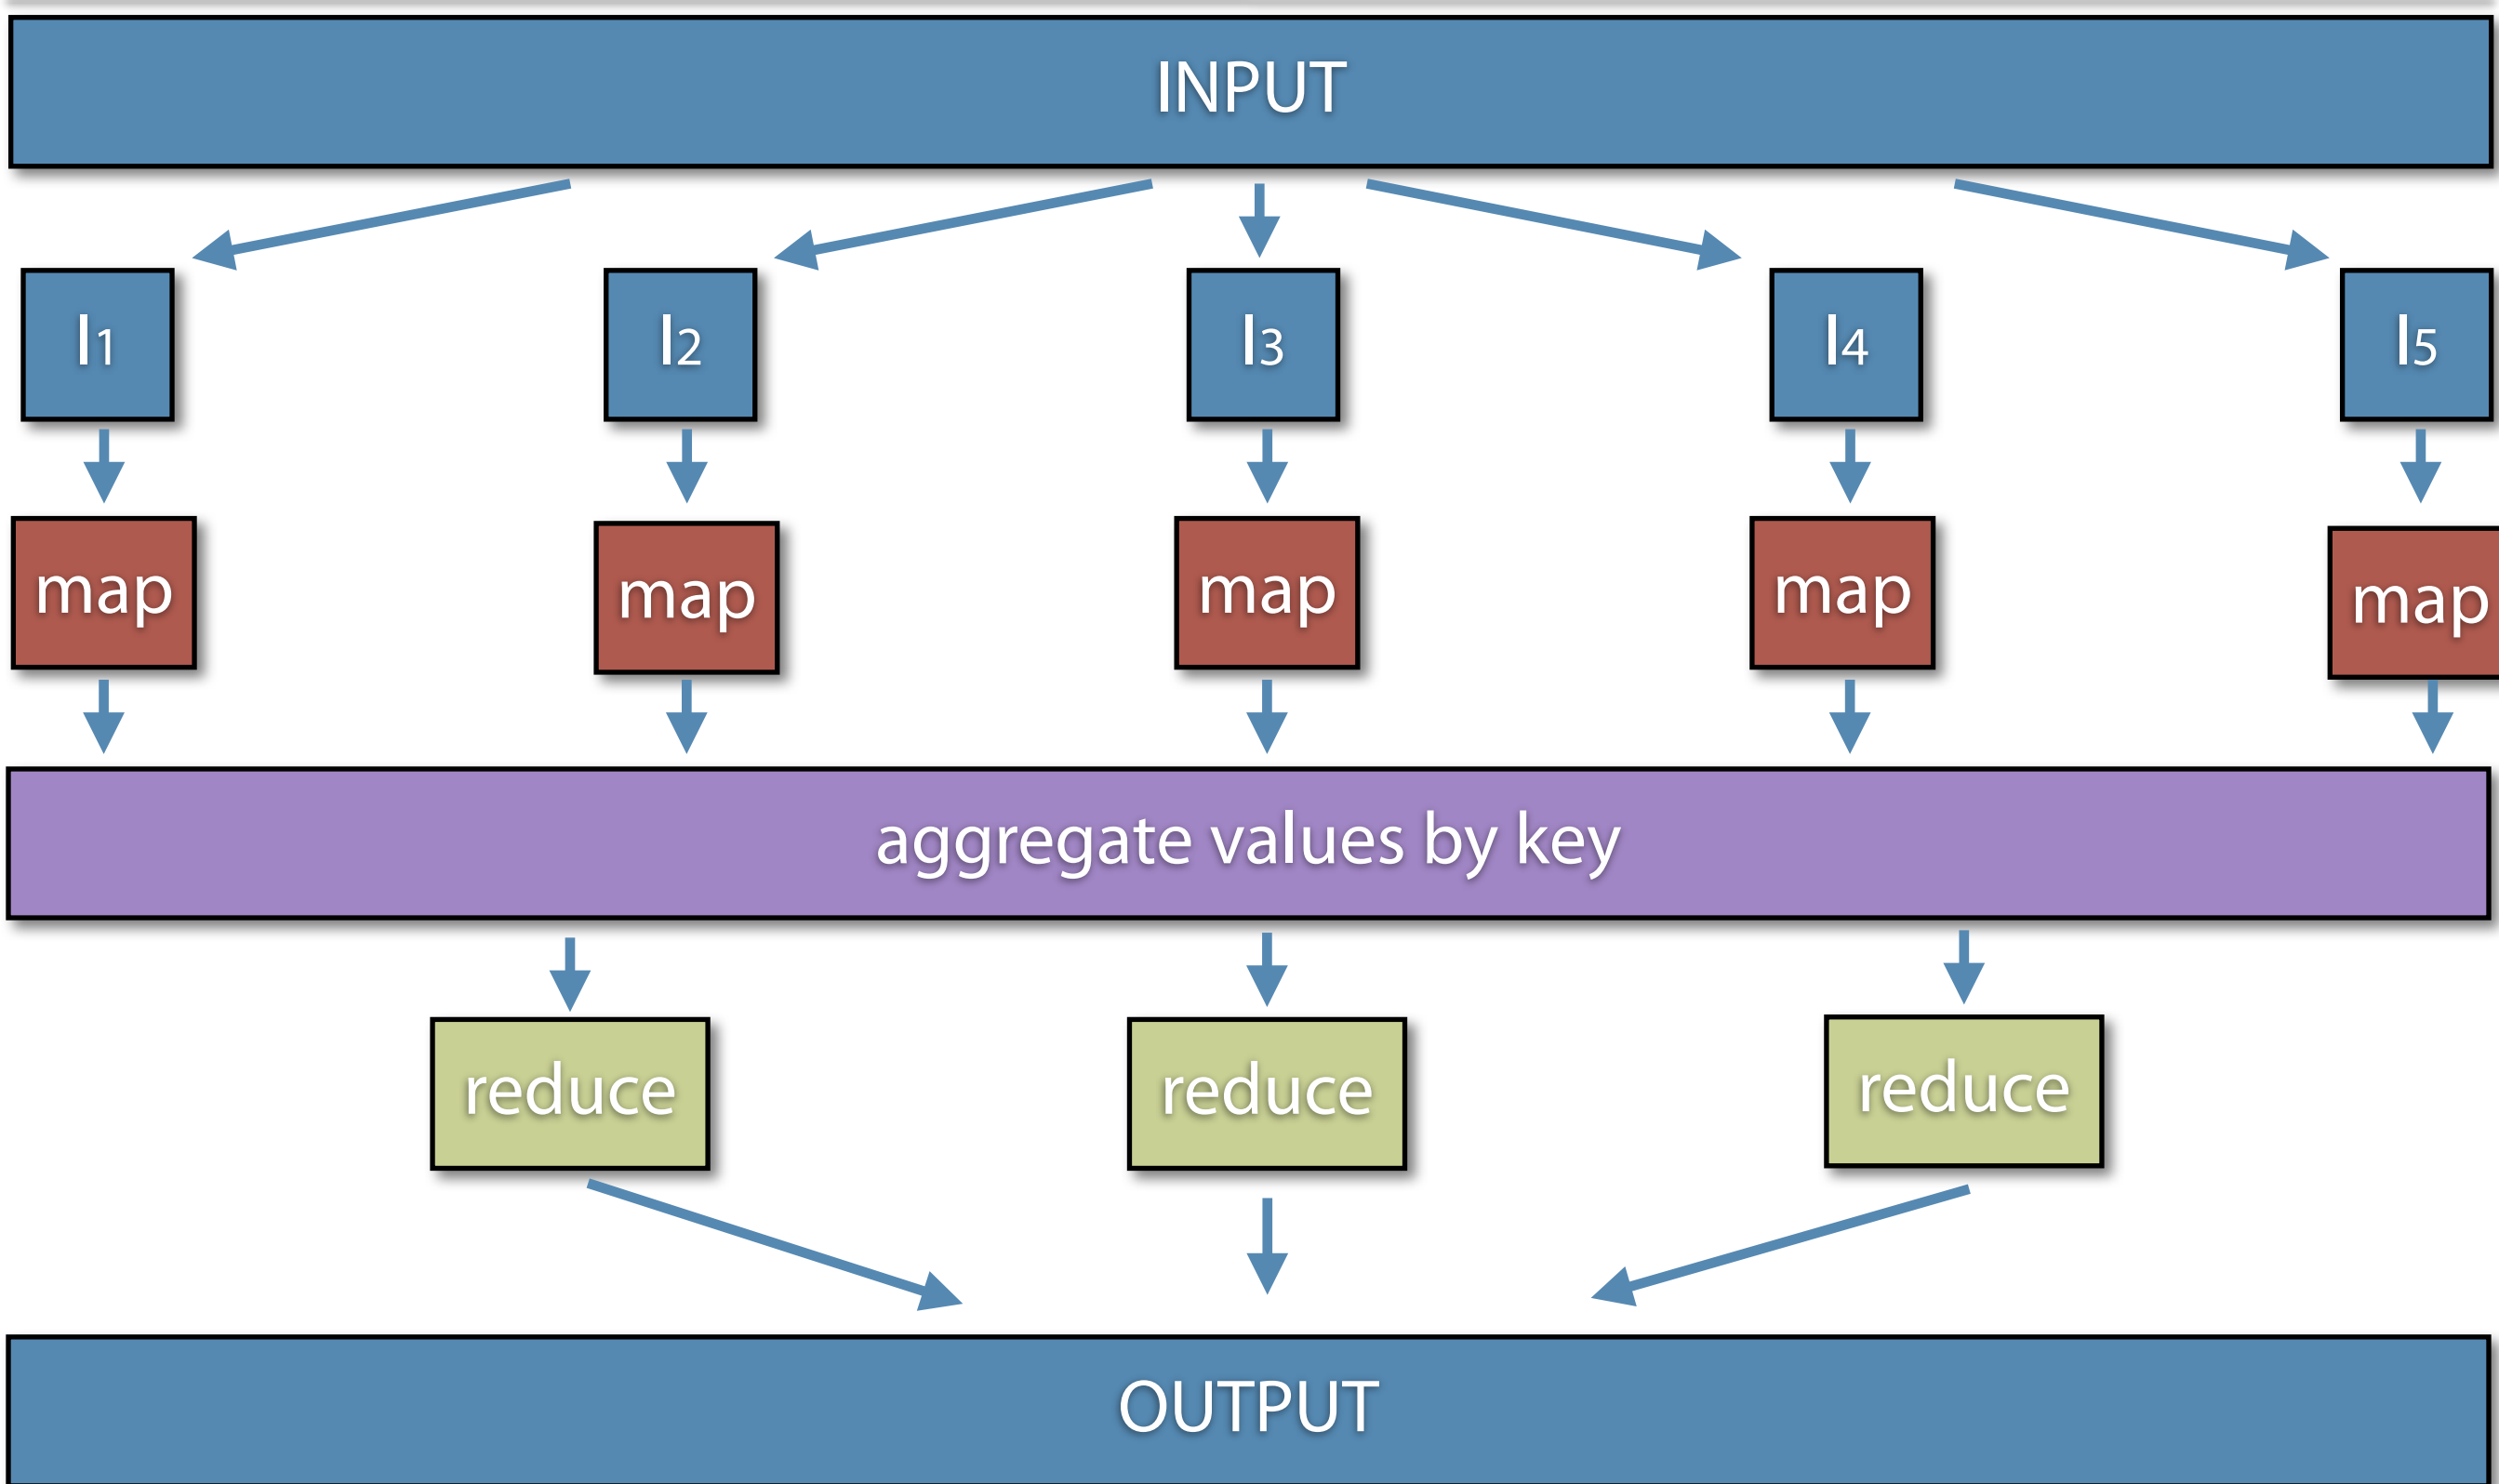
From functional programming...



...to Map Reduce

- **Programmers specify two functions**
 - **map** $(k_1, v_1) \rightarrow [(k_2, v_2)]$
 - **reduce** $(k_2, [v_2]) \rightarrow [(k_3, v_3)]$
- **Map**
 - Receives as input a key-value pair
 - Produces as output a list of key-value pairs
- **Reduce**
 - Receives as input a key-list of values pair
 - Produces as output a list of key-value pairs (typically just one)
- **The runtime support handles everything else...**

Programming Model (simple)

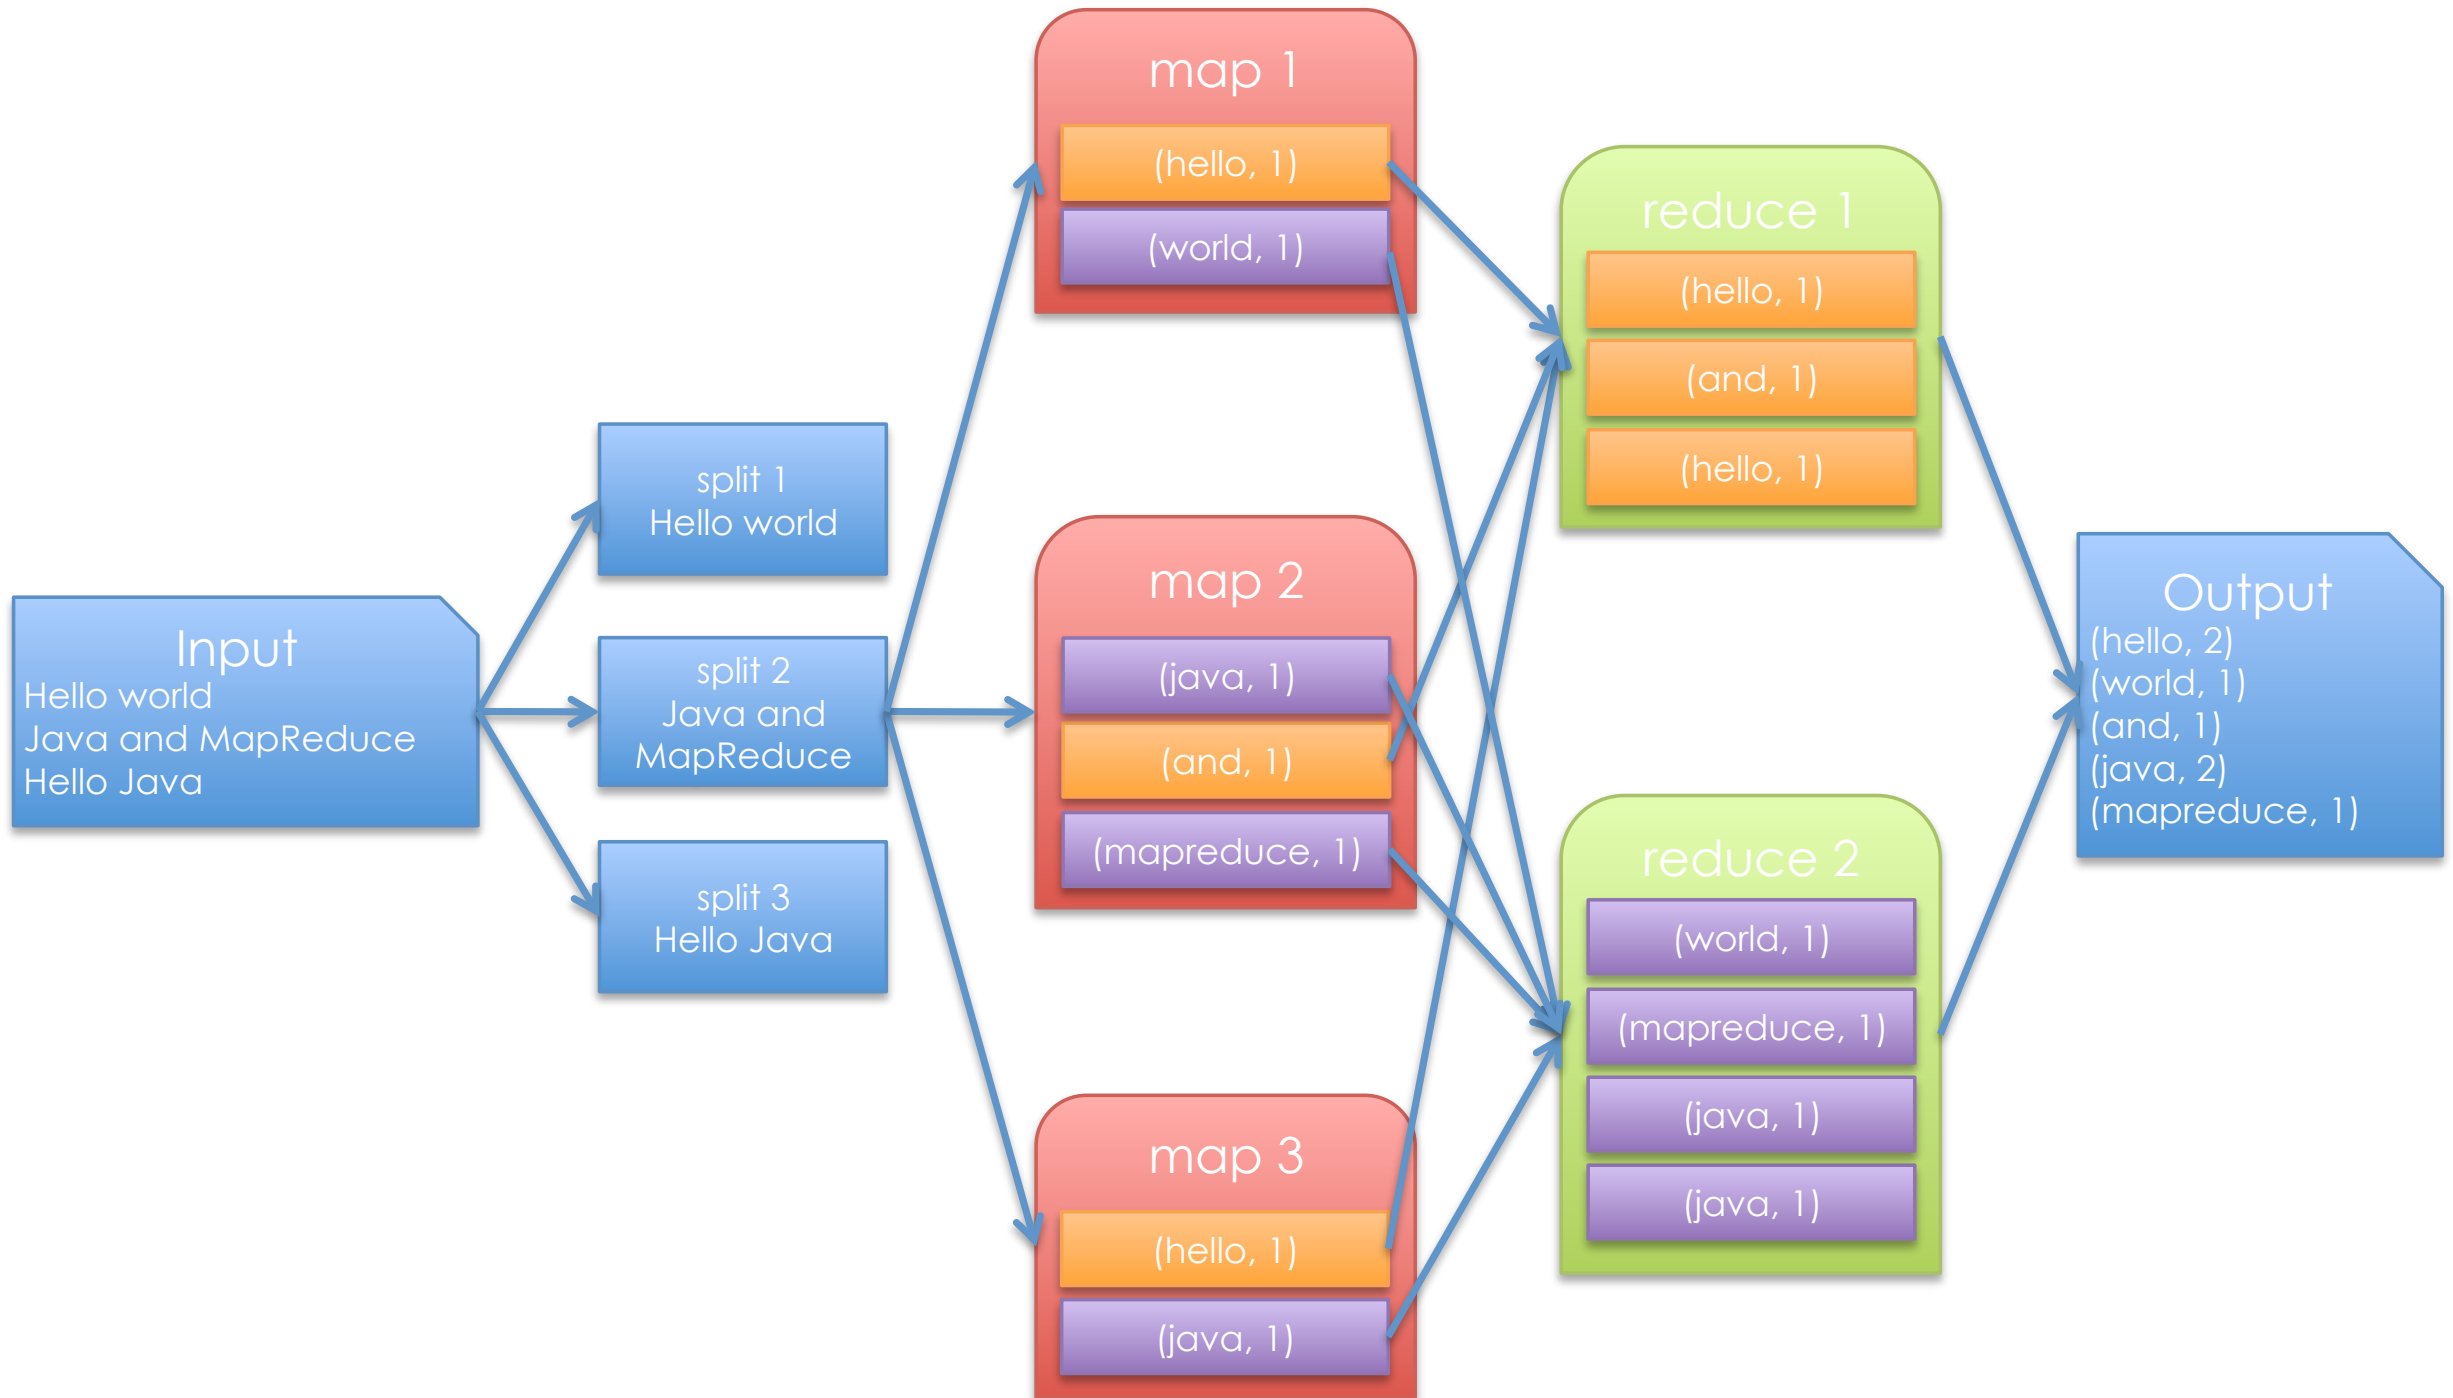


Example (I)

```
1: class MAPPER
2:     method MAP(docid a, doc d)
3:         for all term t  $\in$  doc d do
4:             EMIT(term t, count 1)

1: class REDUCER
2:     method REDUCE(term t, counts [c1, c2, ...])
3:         sum  $\leftarrow$  0
4:         for all count c  $\in$  counts [c1, c2, ...] do
5:             sum  $\leftarrow$  sum + c
6:         EMIT(term t, count sum)
```

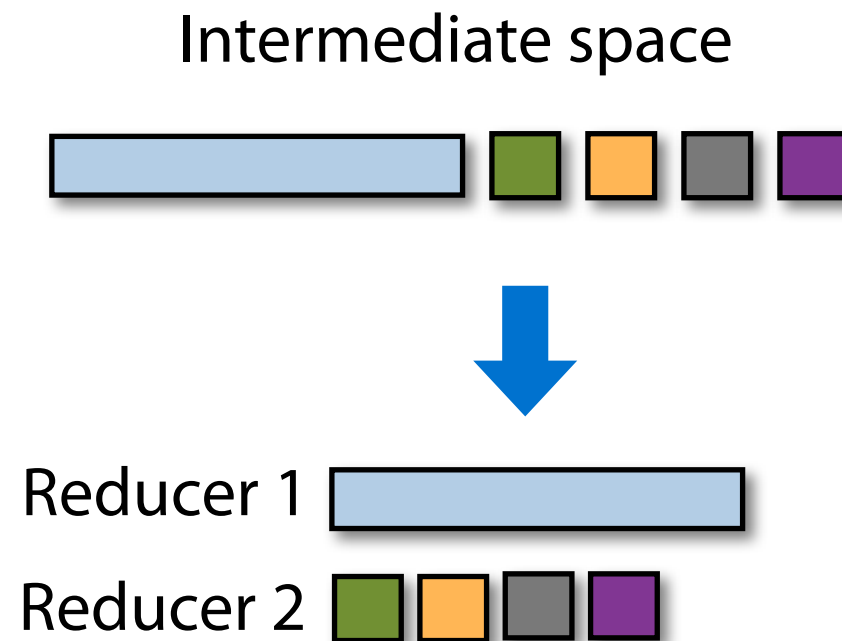

Example (II)



- **Handles scheduling**
 - Assigns workers to map and reduce tasks
- **Handles “data distribution”**
 - Moves processes to data
- **Handles synchronization**
 - Gathers, sorts, and shuffles intermediate data
- **Handles errors and faults**
 - Detects worker failures and restarts
- **Everything happens on top of a distributed FS**

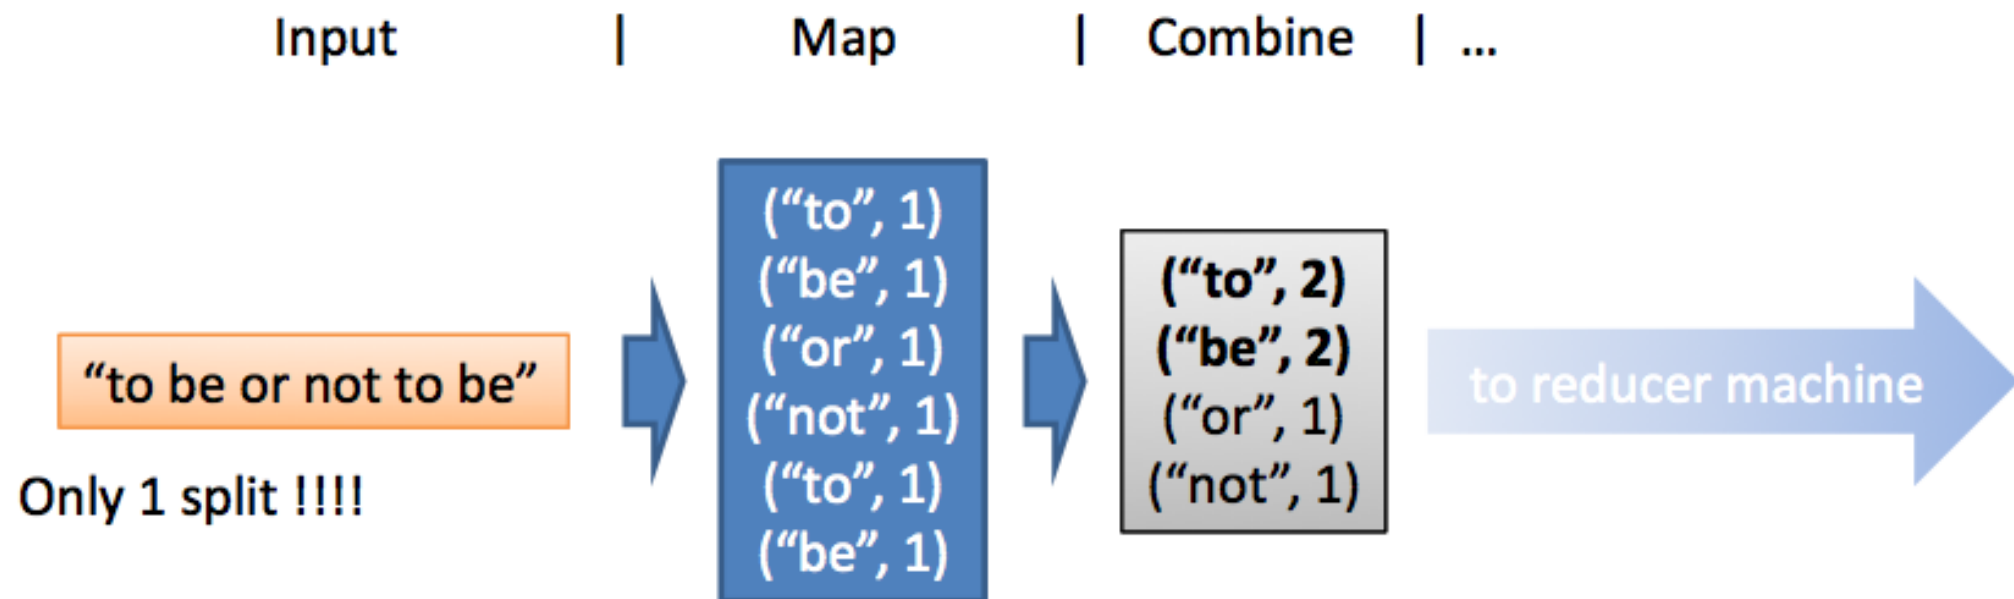
Partitioners

- **Balance the key assignments to reducers**
 - By default, intermediate keys are hashed to reducers
 - Partitioner specifies the node to which an intermediate key-value pair must be copied
 - Divides up key space for parallel reduce operations
 - Partitioner only considers the key and ignores the value

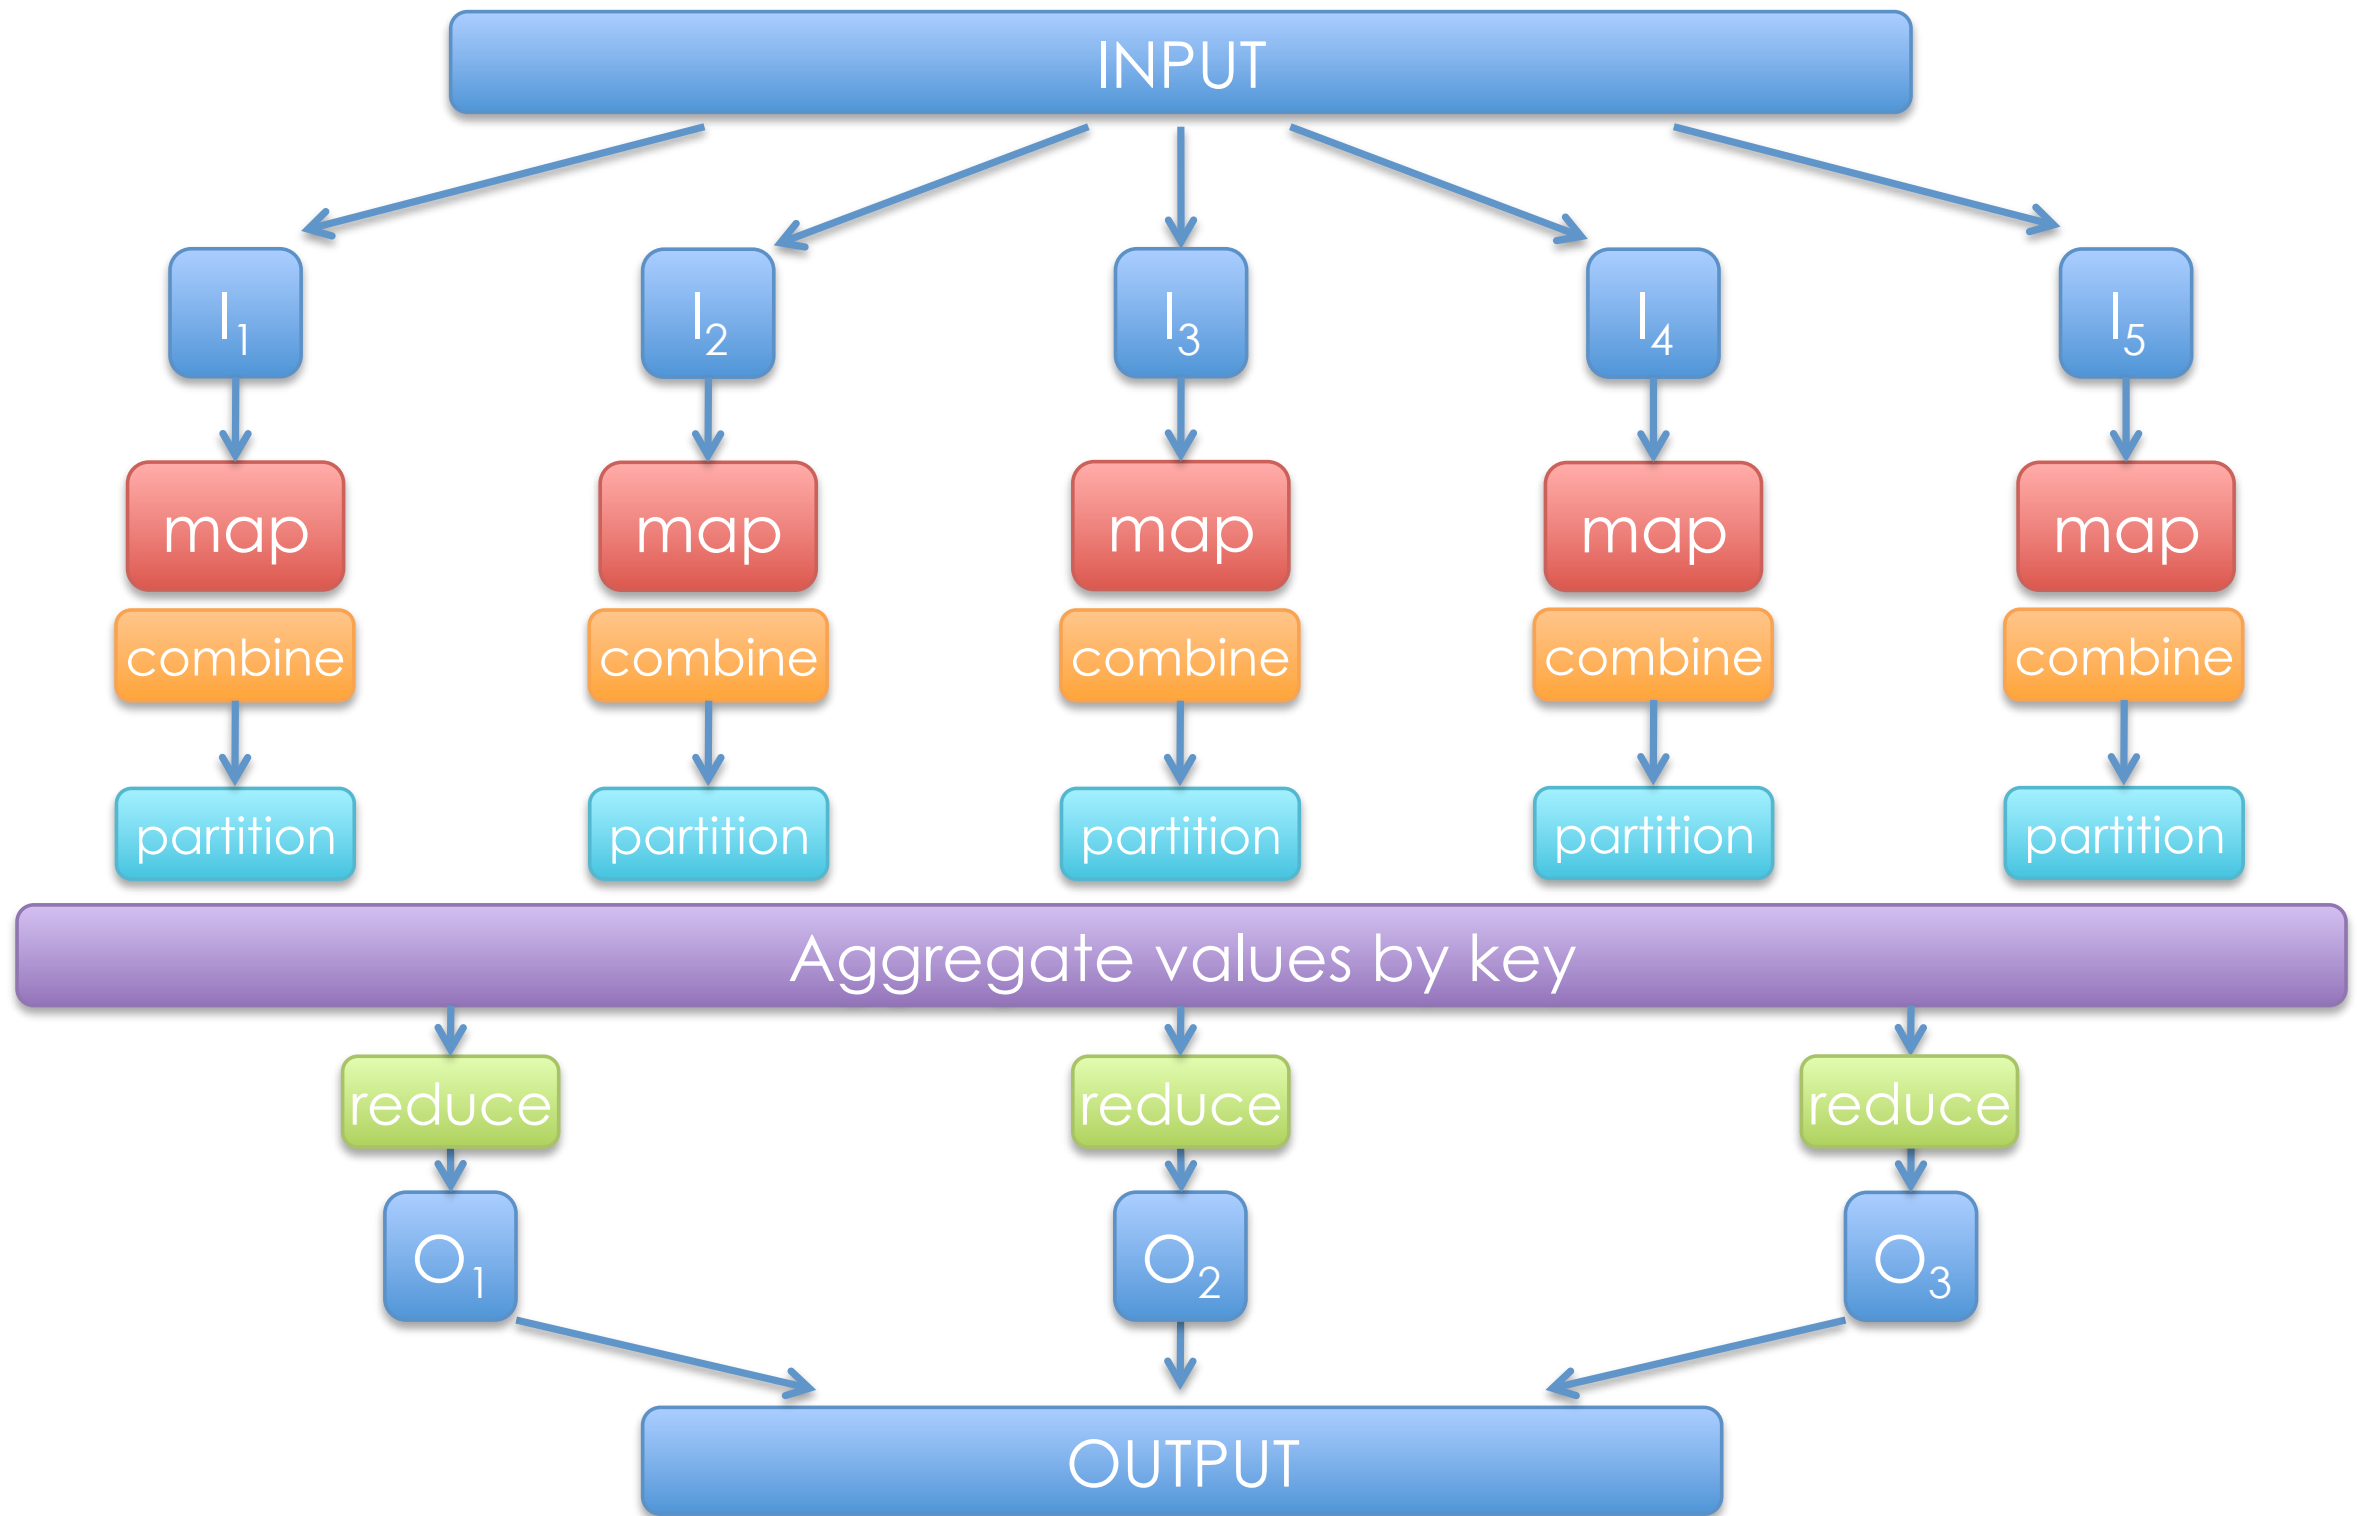


- **Local aggregation before the shuffle**

- All the key-value pairs from mappers need to be copied across the network
- The amount of intermediate data may be larger than the input collection itself
- Perform local aggregation on the output of each mapper (same machine)
- Typically, a combiner is a (local) copy of the reducer



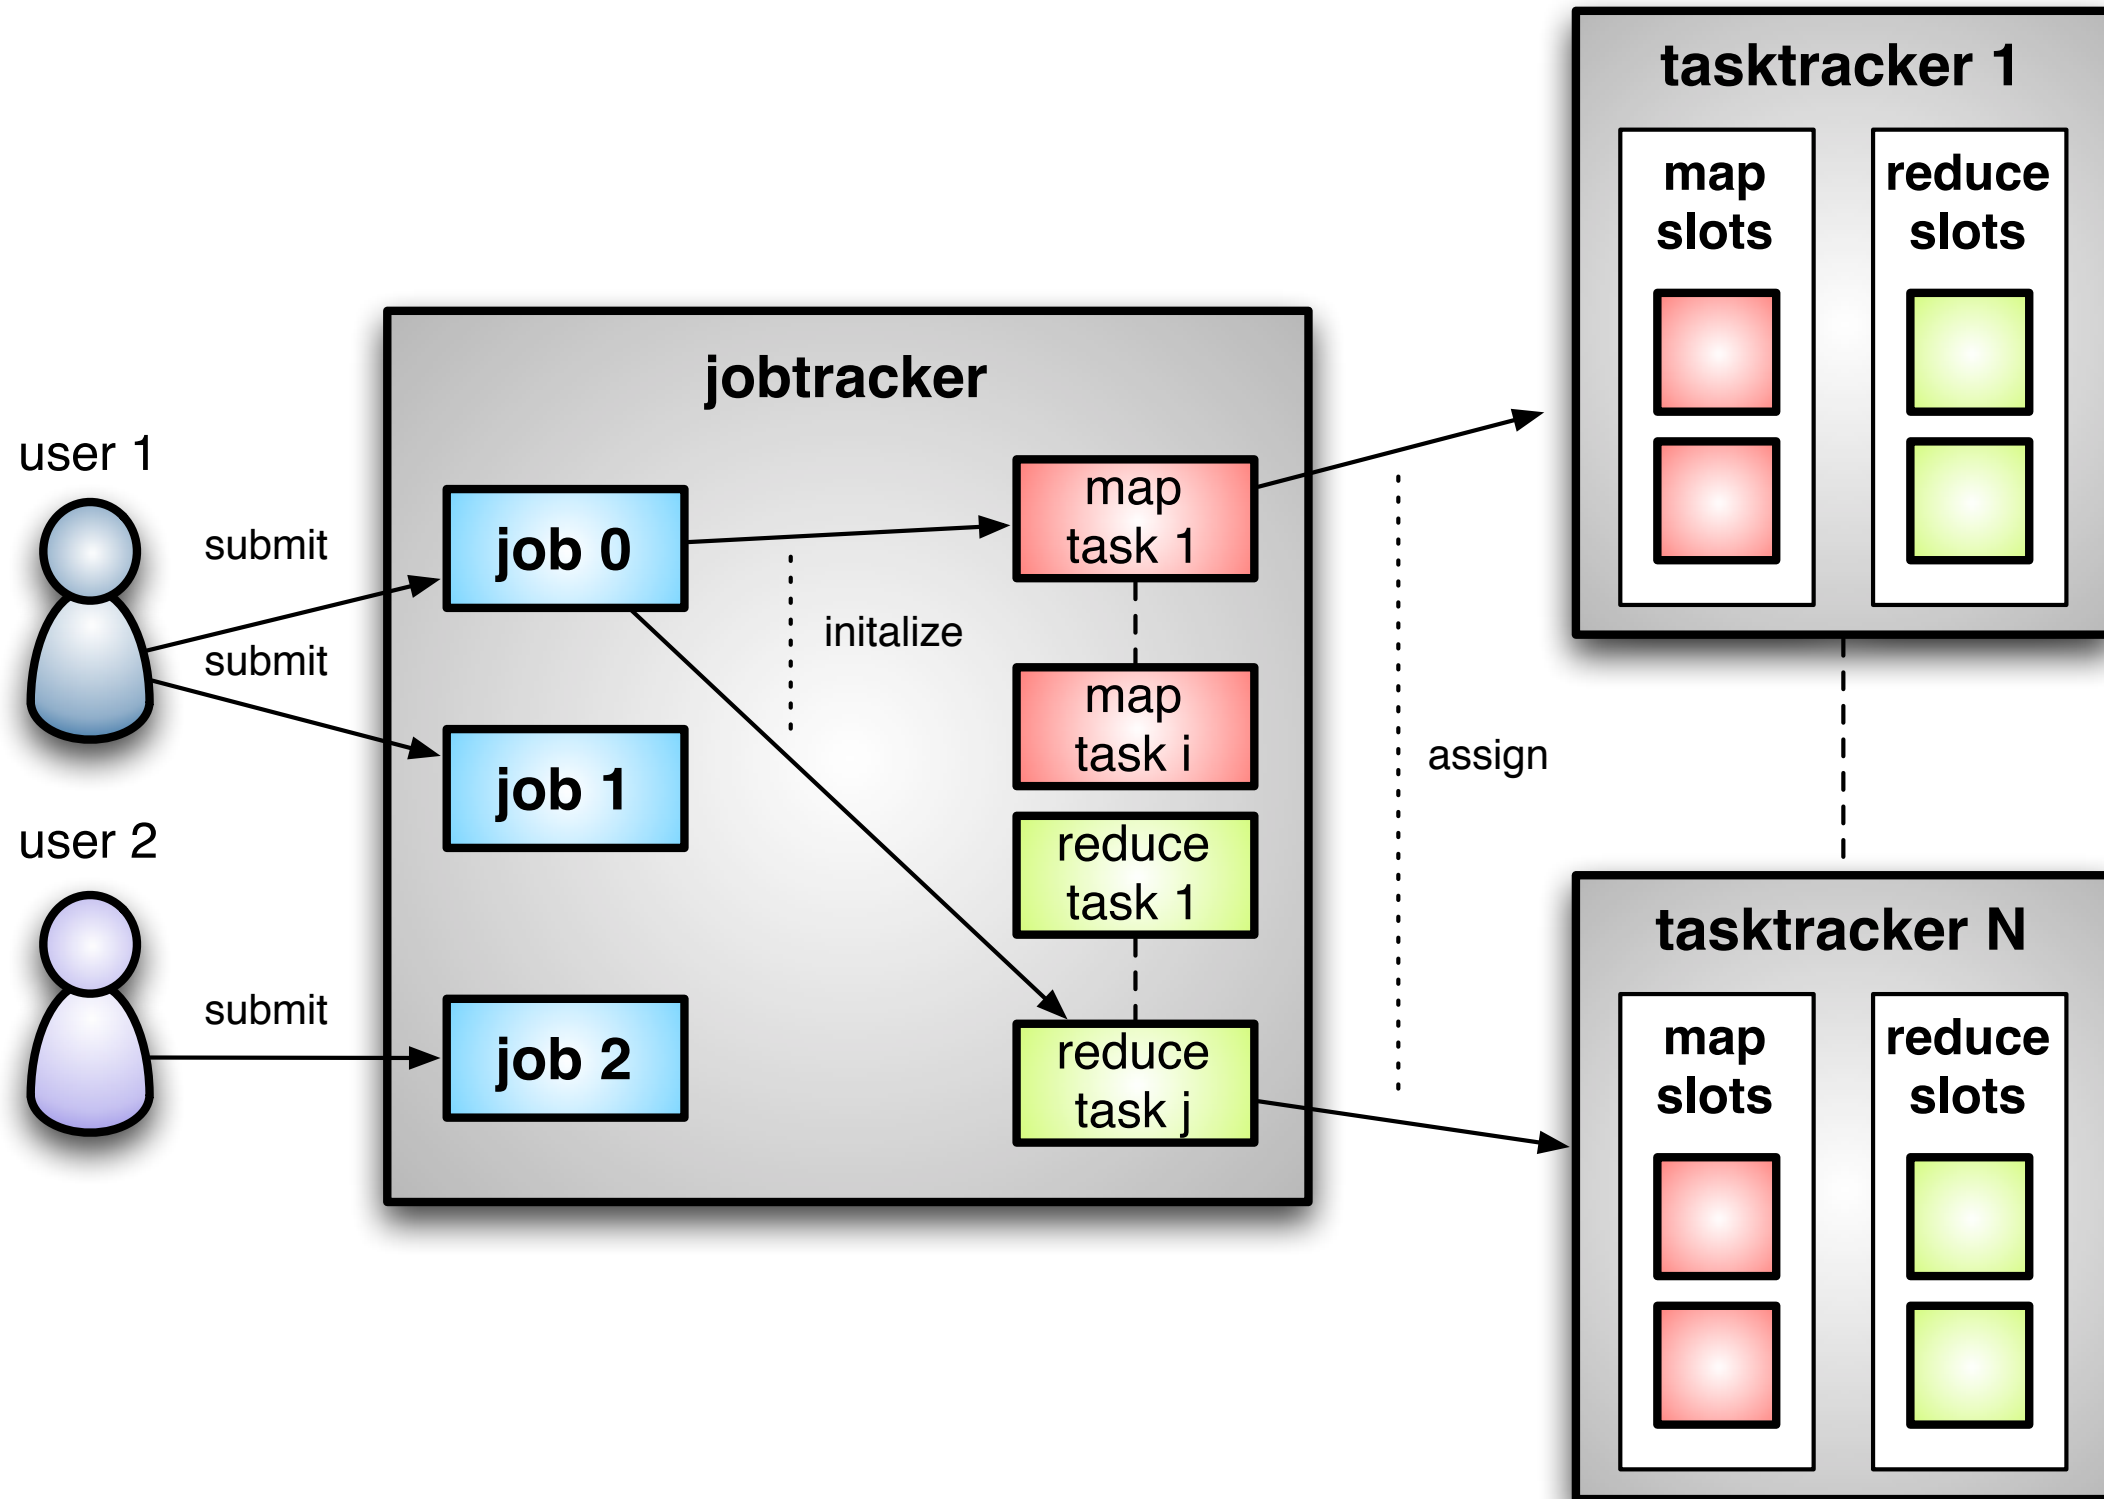
Programming Model (complete)



Terminology

- **Job**
- **Task**
- **Slot**
- **JobTracker**
 - Accepts Map/Reduce jobs submitted by users
 - Assigns Map and Reduce tasks to Task Trackers
 - Monitors task and Task Tracker status, re-executes tasks upon failure
- **TaskTracker**
 - Run Map and Reduce tasks upon instruction from the Job Tracker
 - Manage storage and transmission of intermediate output
- **Splits**
 - Data locality optimization

Runtime



Scheduling

- **One master, many workers**

- Input data split into M map tasks (typically 64 MB in size)
- Reduce phase partitioned into R reduce tasks ($\text{hash}(k) \bmod R$)
- Tasks are assigned to workers dynamically
- Often: $M=200,000$; $R=4000$; workers=2000

- **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

Parallelism

- **Map functions run in parallel, create intermediate values from each input data set**
 - The programmer must specify a proper input split (chunk) between mappers to enable parallelism
- **Reduce functions also run in parallel, each will work on different output keys**
 - Number of reducers is a key parameter which determines map-reduce performance

Speculative Execution

- **Problem: Stragglers (i.e., slow workers) significantly lengthen the completion time**
 - Other jobs may be consuming resources on machine
 - Bad disks with soft (i.e., correctable) errors transfer data very slowly
 - Other weird things: processor caches disabled at machine init
- **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.**

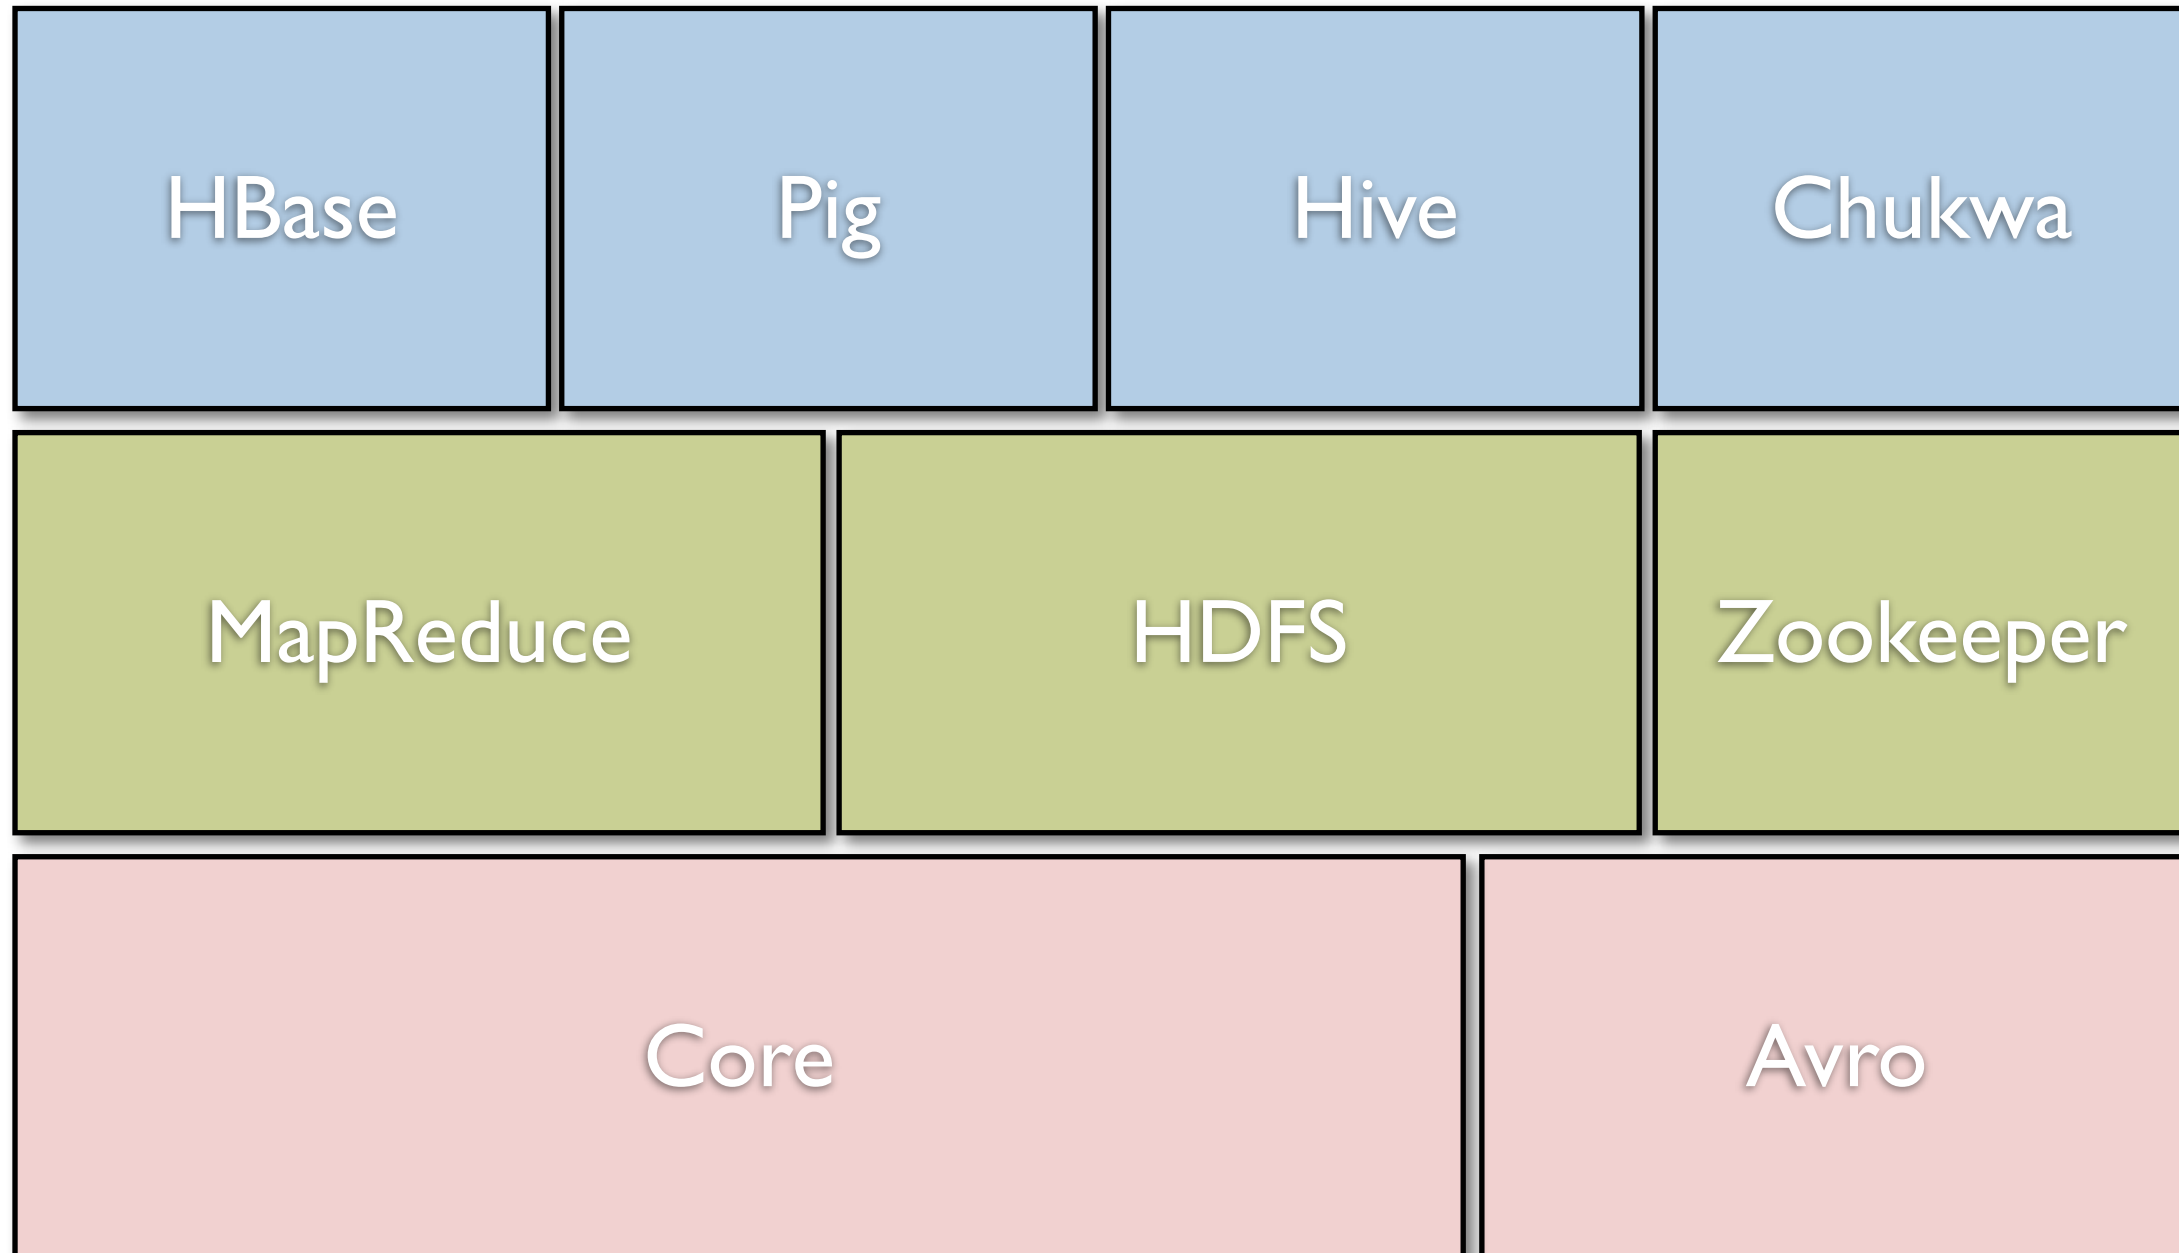
Fault Tolerance

- **Master keeps track of progress of each task and worker nodes**
 - If a node fails, it re-executes the completed as well as in-progress map tasks on other nodes that are alive
 - It also executes in-progress reduce tasks.
- **If particular input key/value pairs keep crashing**
 - Master blacklists them and skips them from re-execution
- **Tolerate small failures, allow the job to run in best-effort basis**
 - For large datasets containing potentially millions of records, we don't want to stop computation for a few records not processing correctly
 - User can set the failure tolerance level

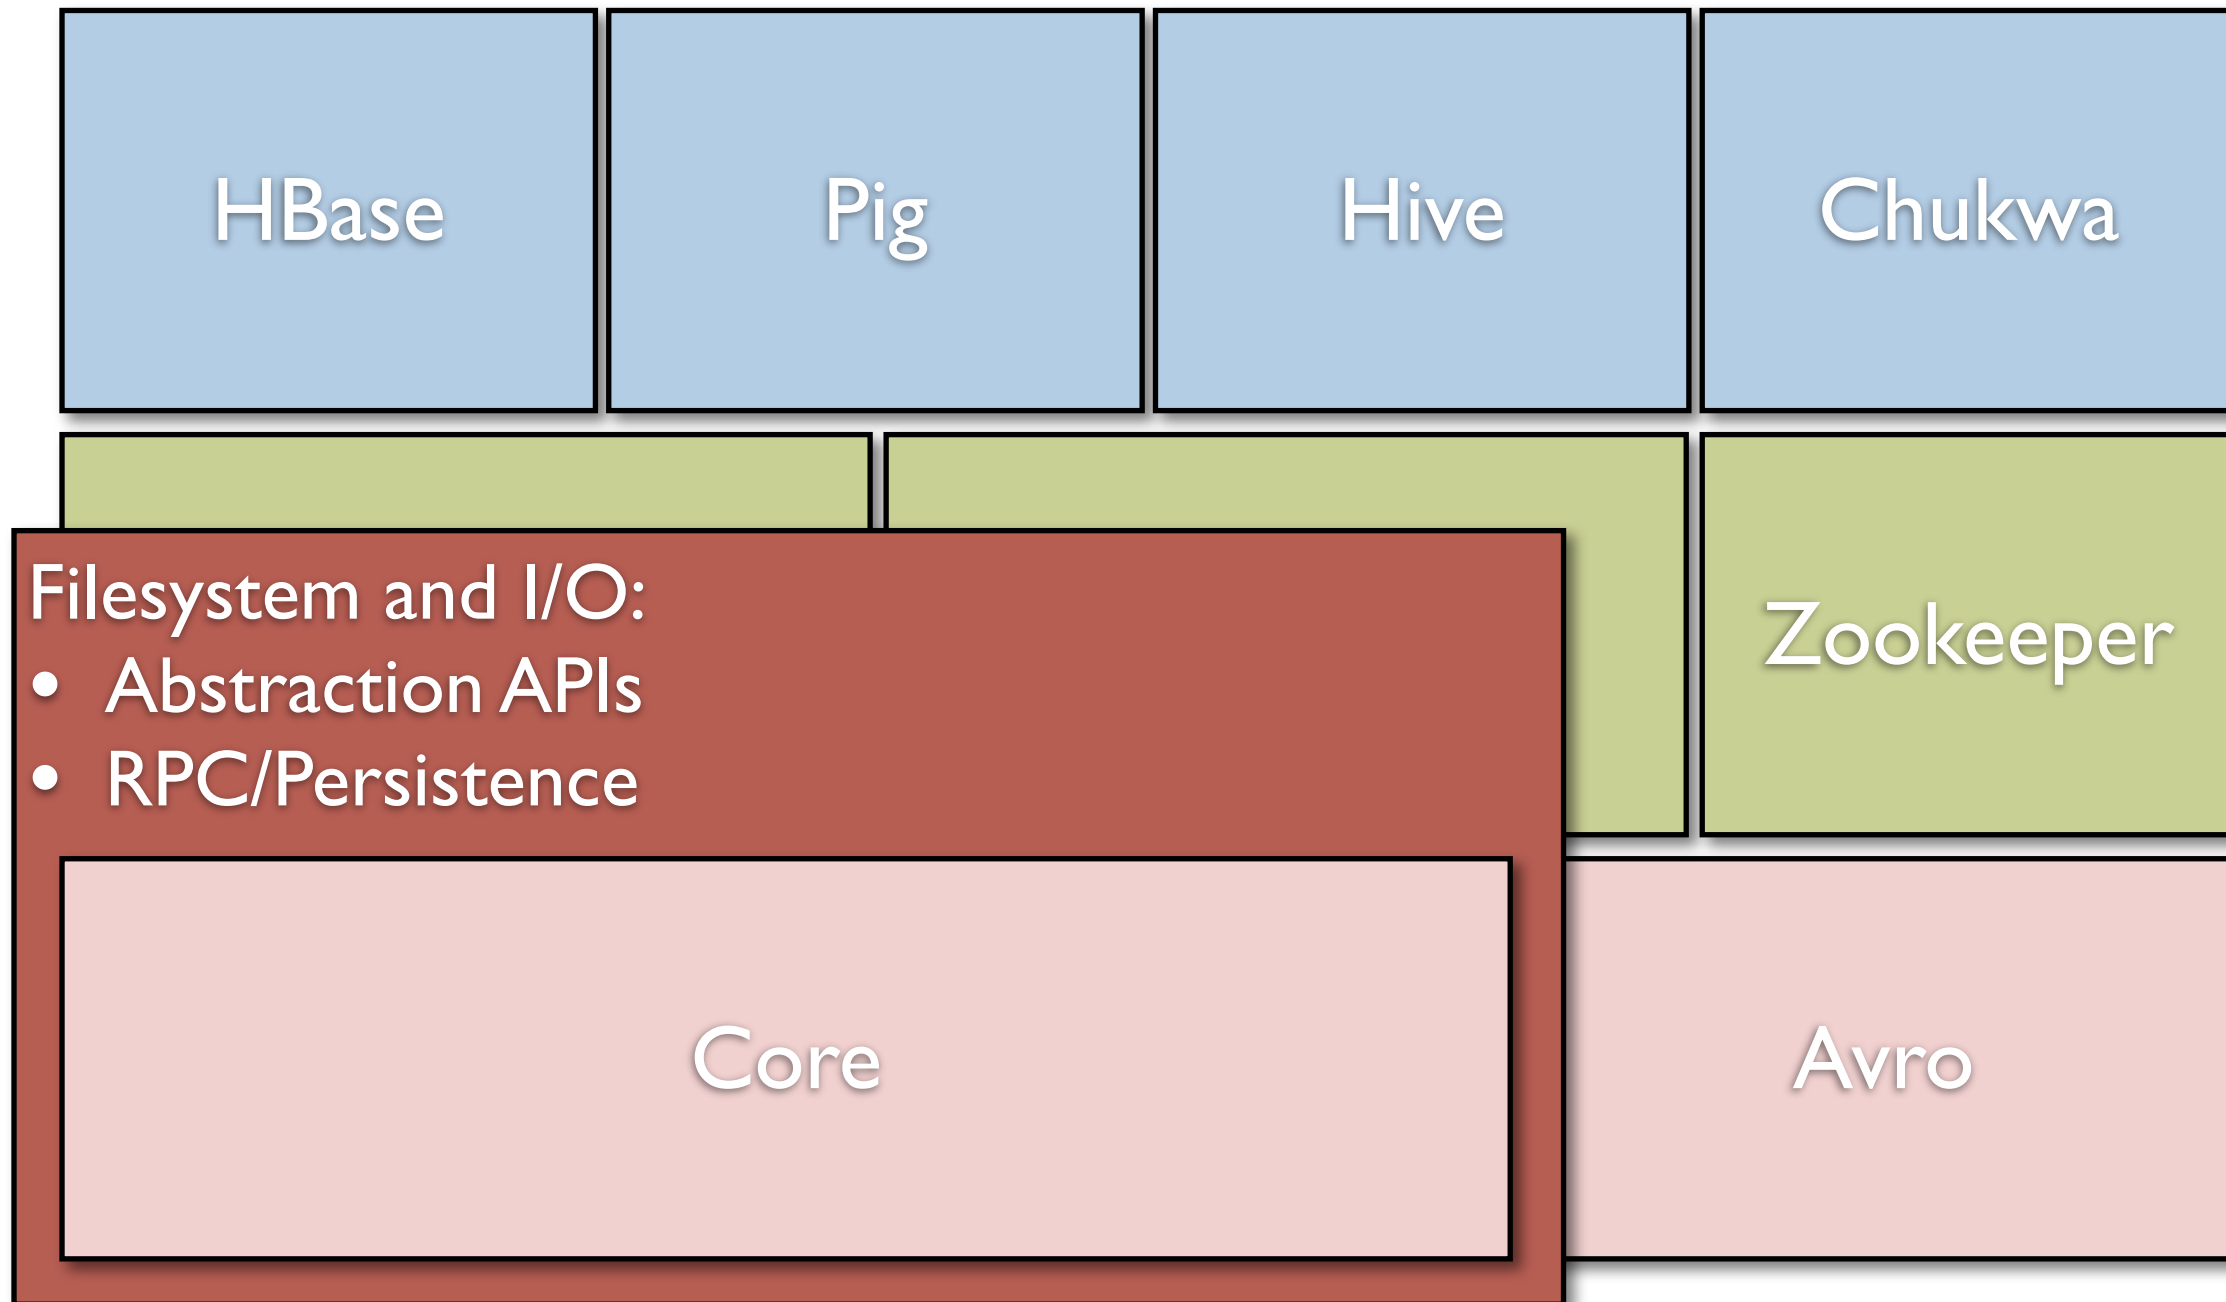
Performance

- **Maximizing Map input transfer rate**
 - Input Locality
 - Minimal deserialization overhead
- **Small intermediate output**
 - $M \times R$ transfers over the network
 - Minimize/compress transfers
 - Avoid shuffling/sorting if possible (e.g. map-only computations)
 - Use combiners and/or partitioners!!!
 - Compress everything (automatic)
- **Opportunity to Load Balance**
- **Changing algorithm to suit architecture yields best implementation**

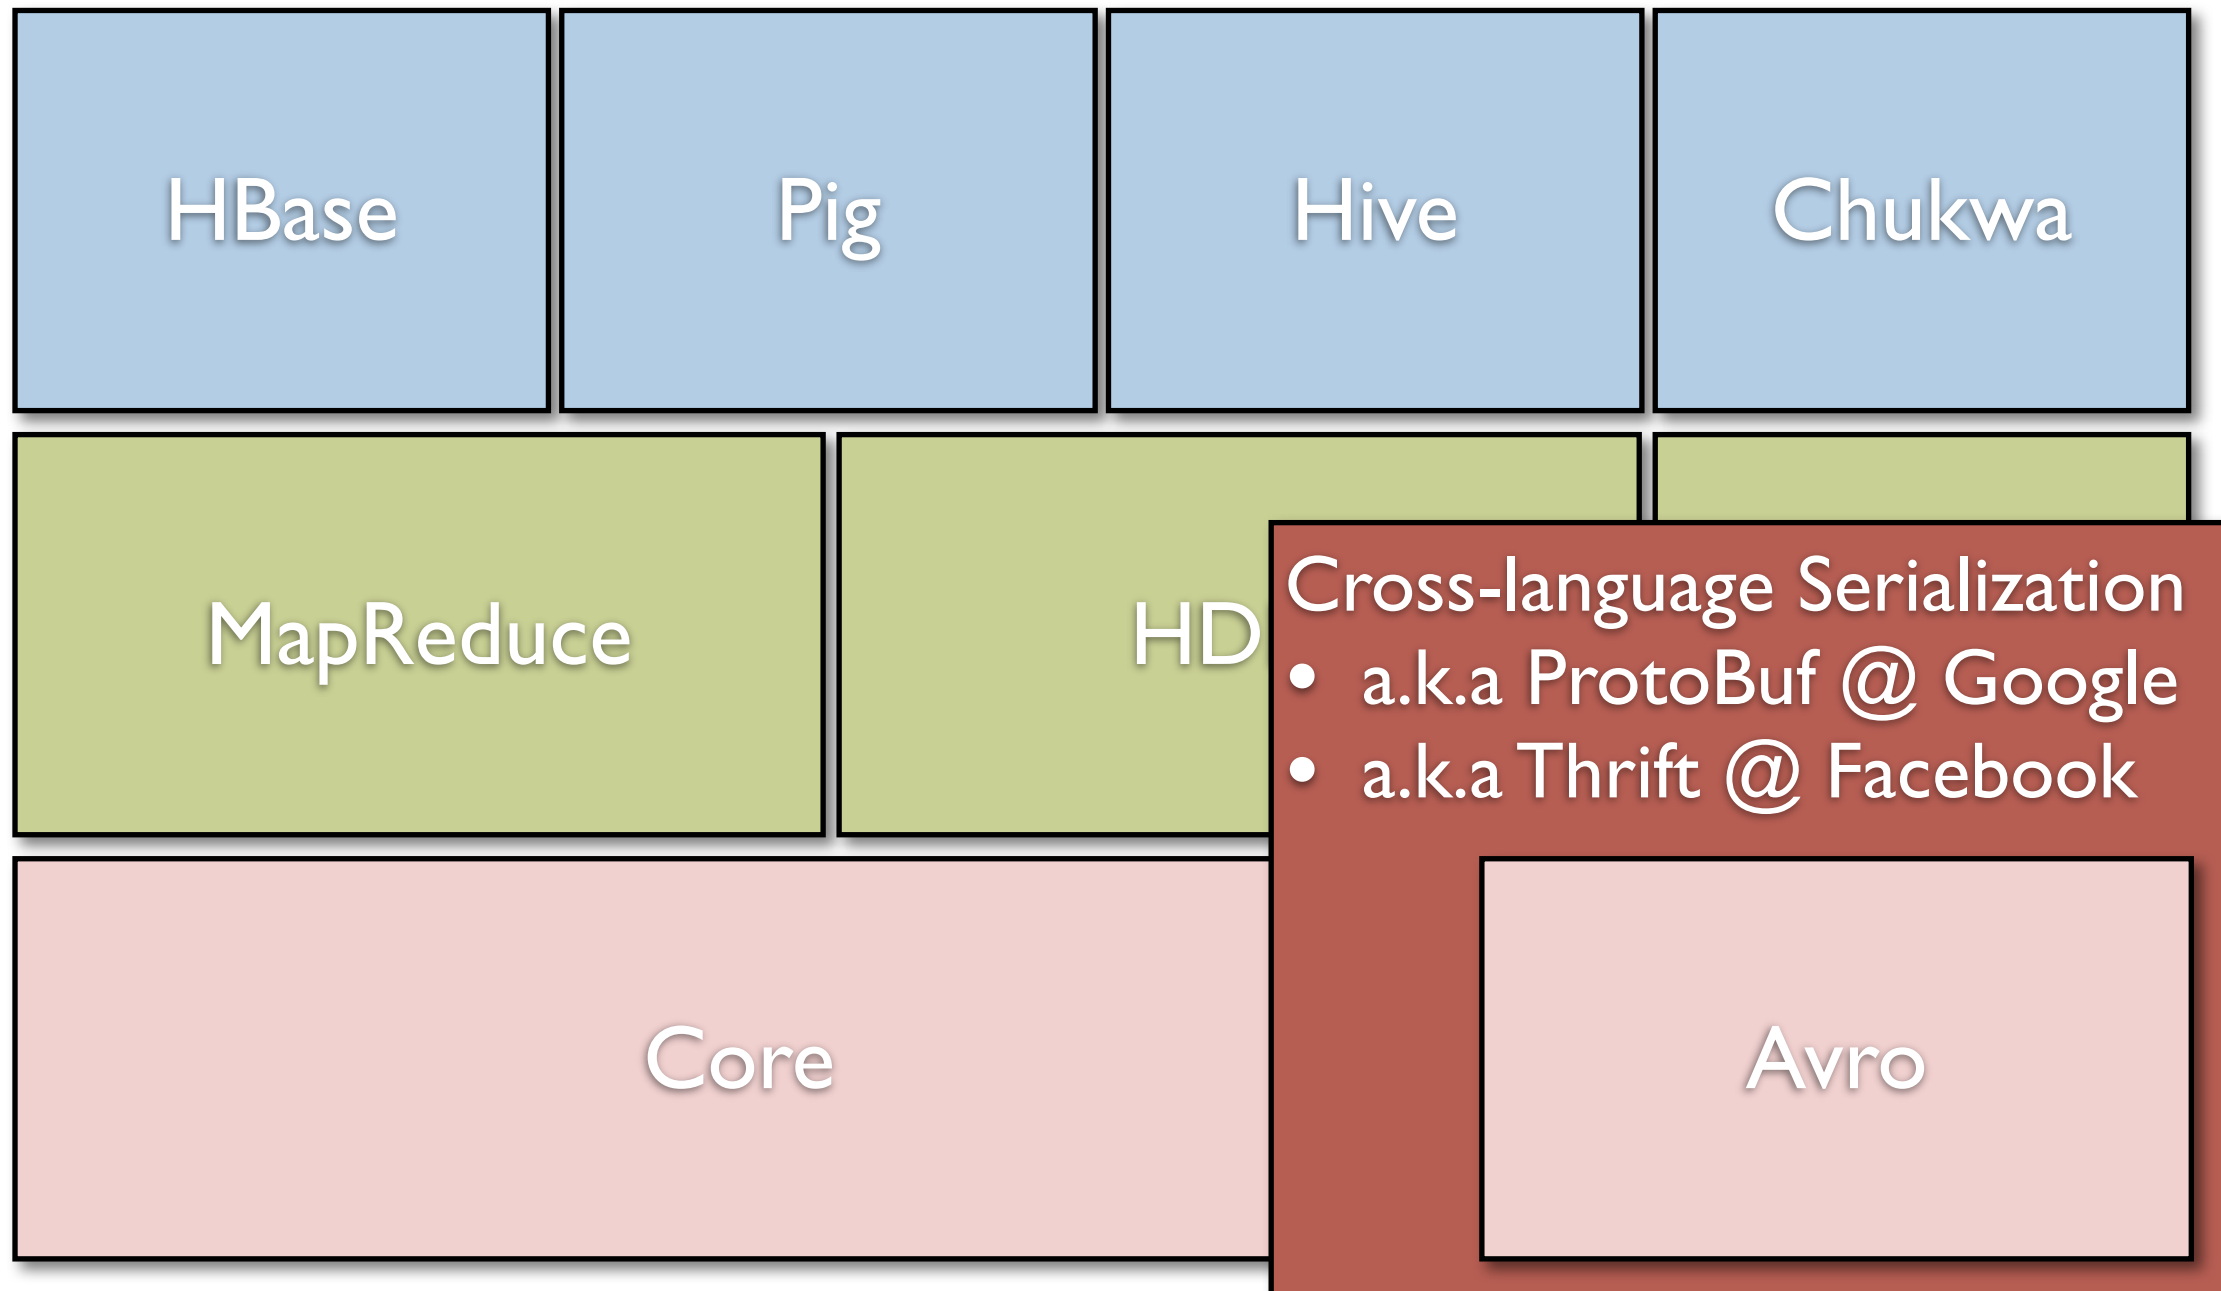
HADOOP



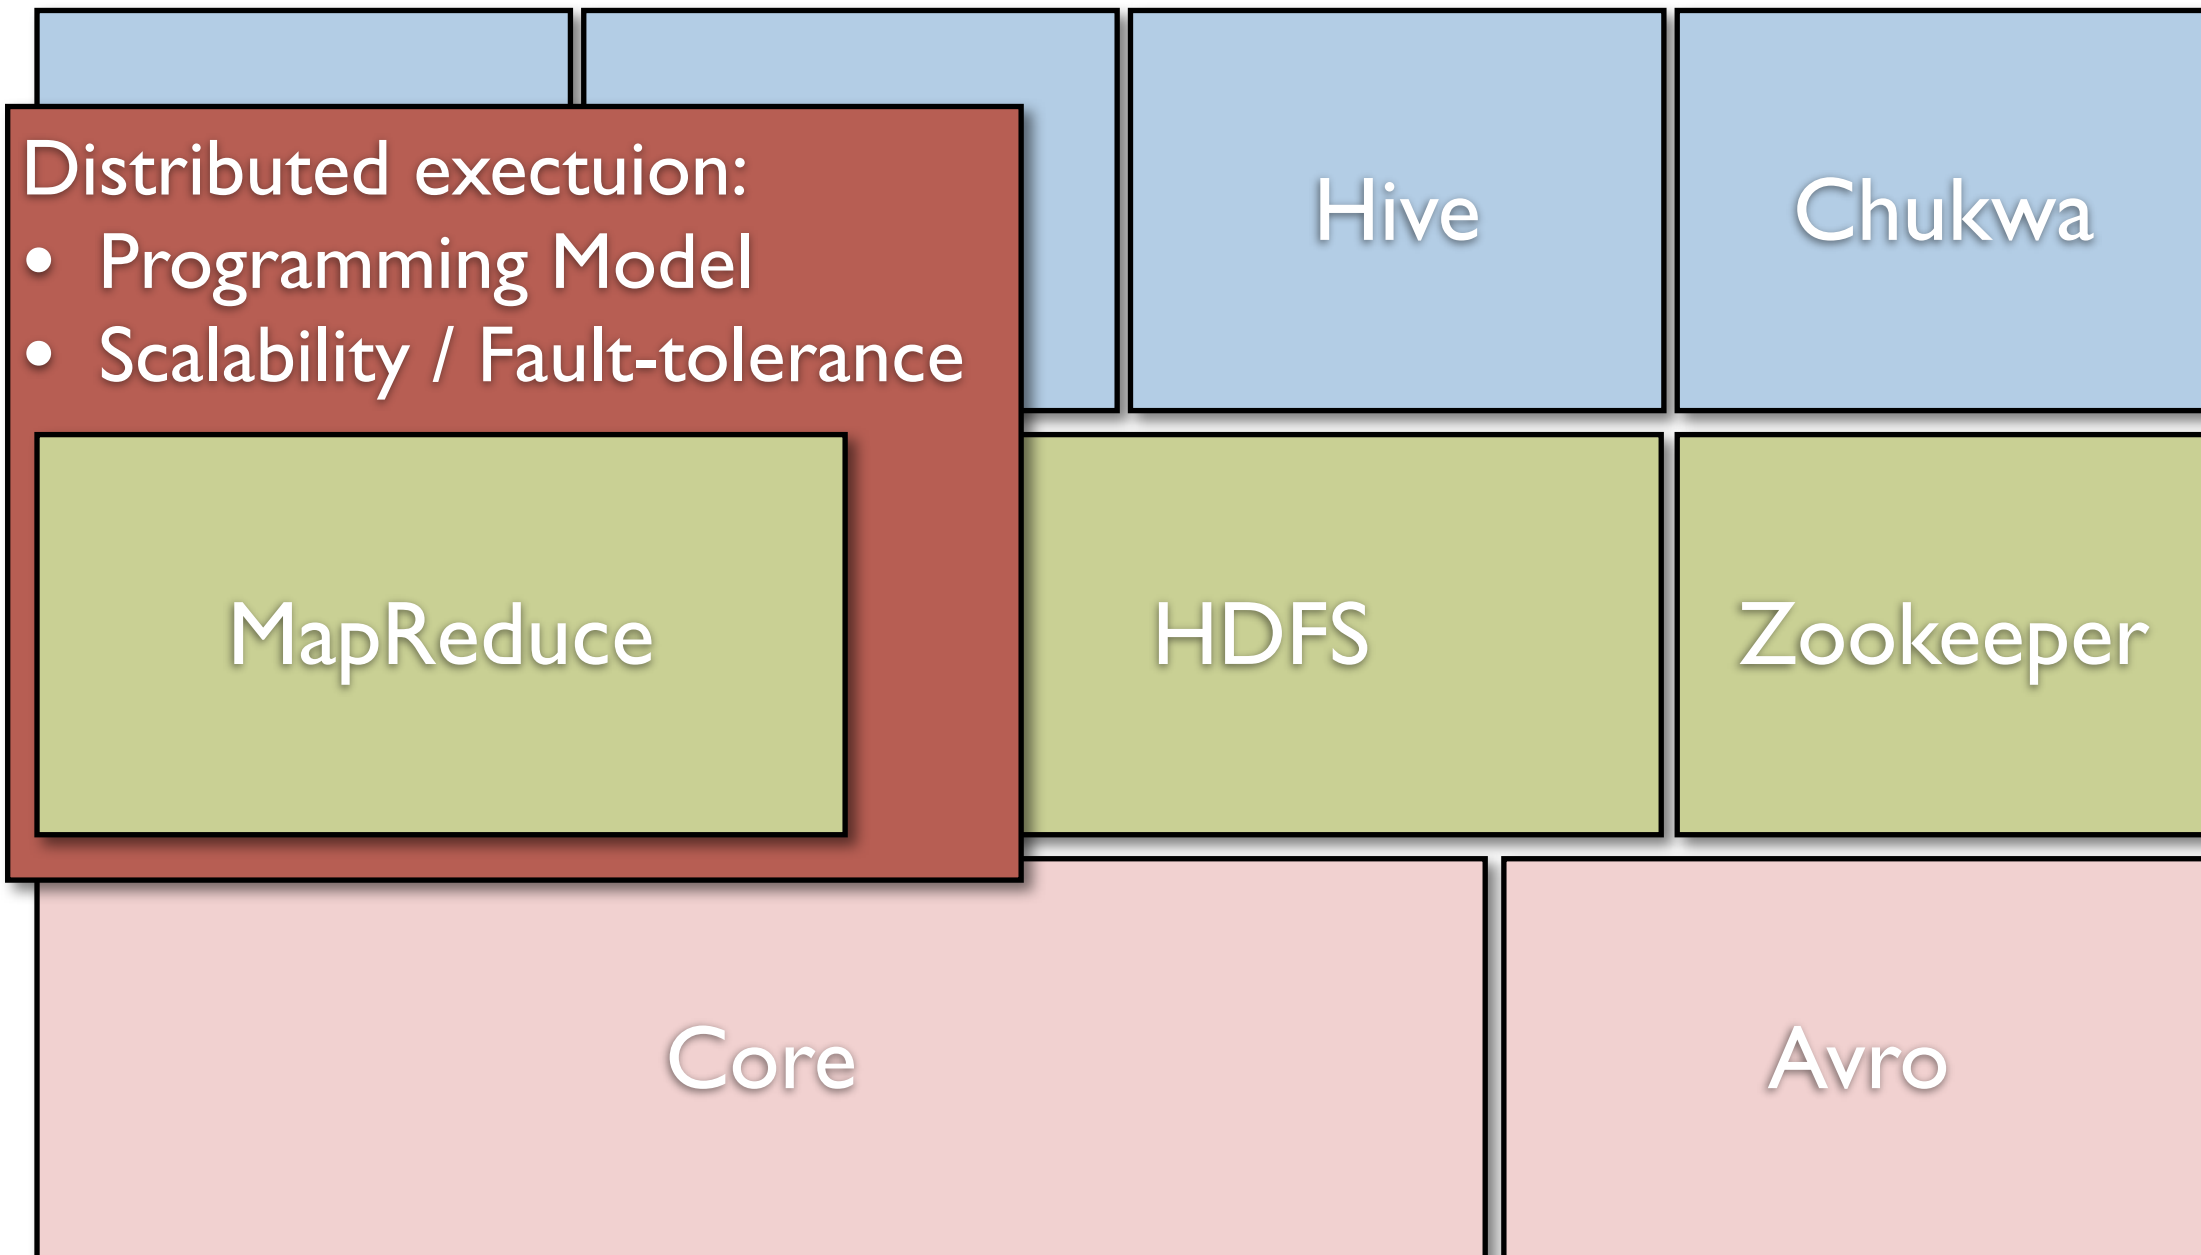
HADOOP



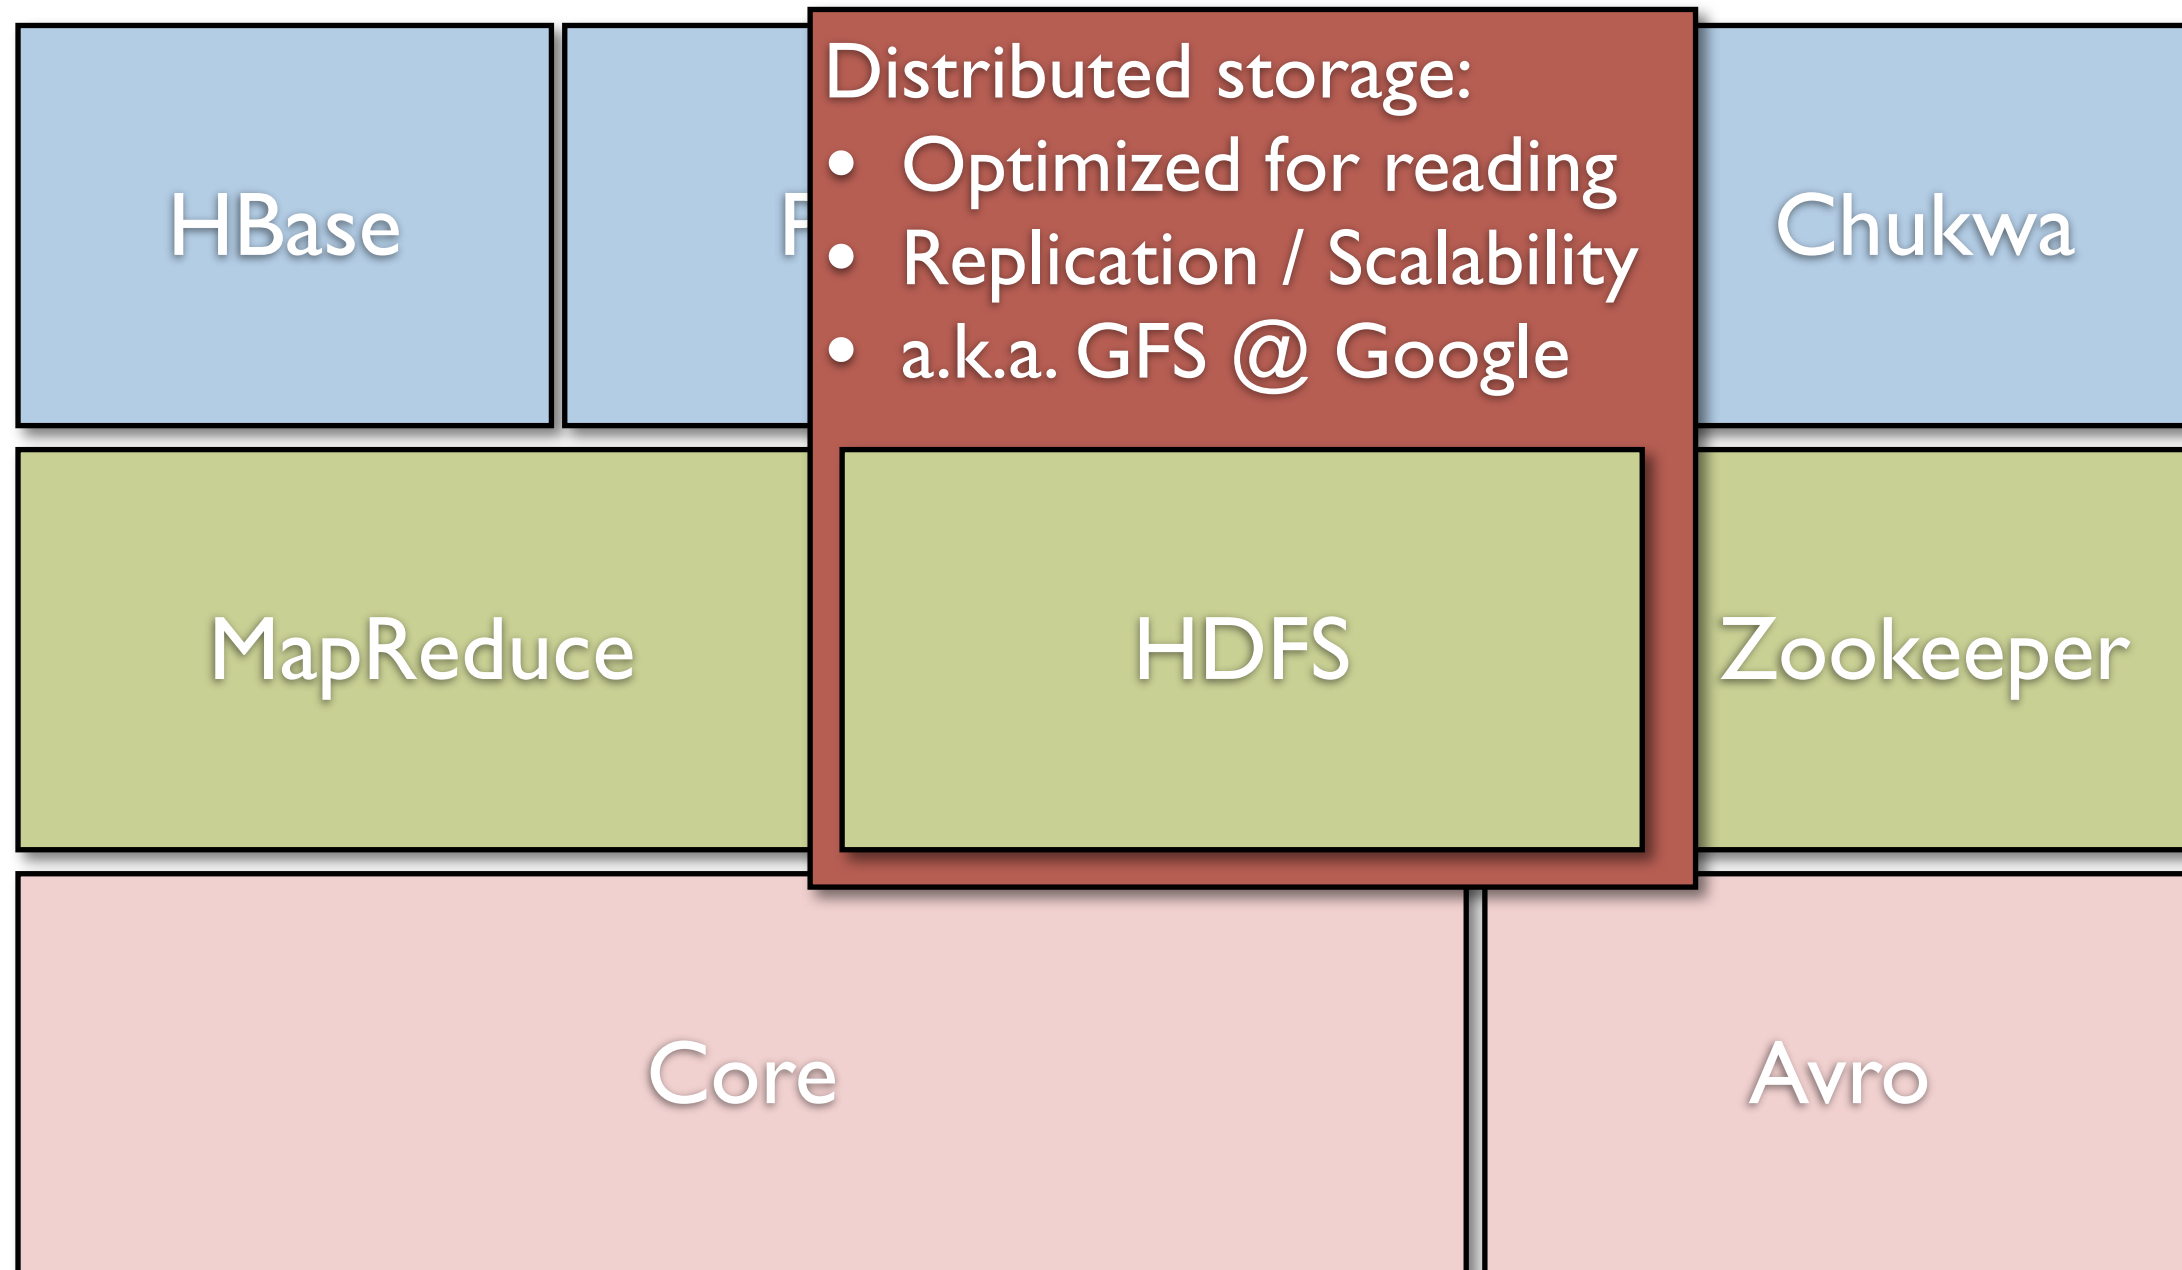
HADOOP

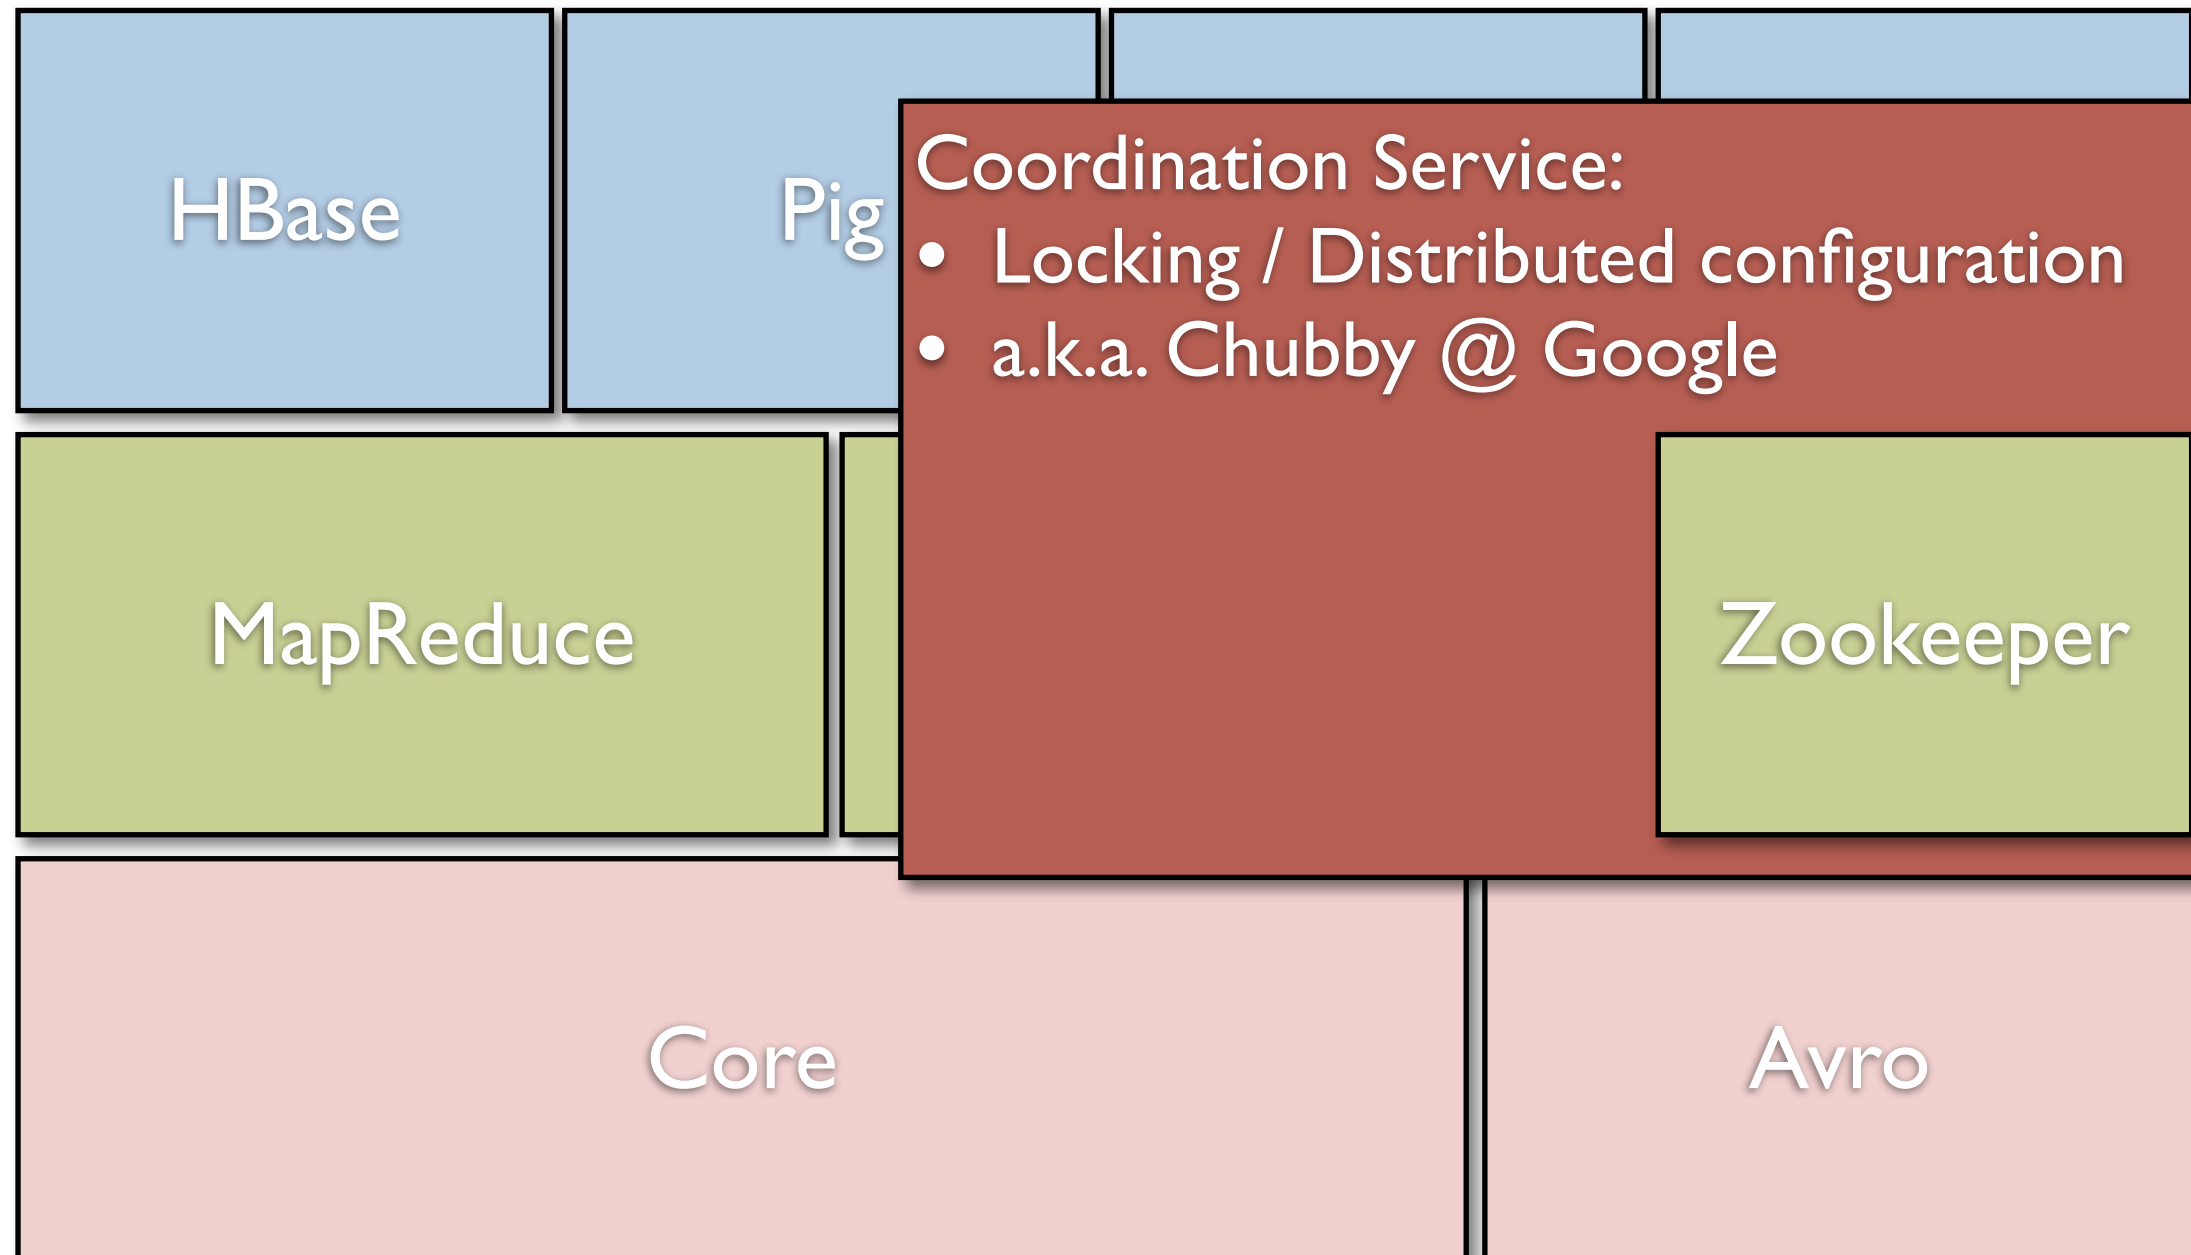


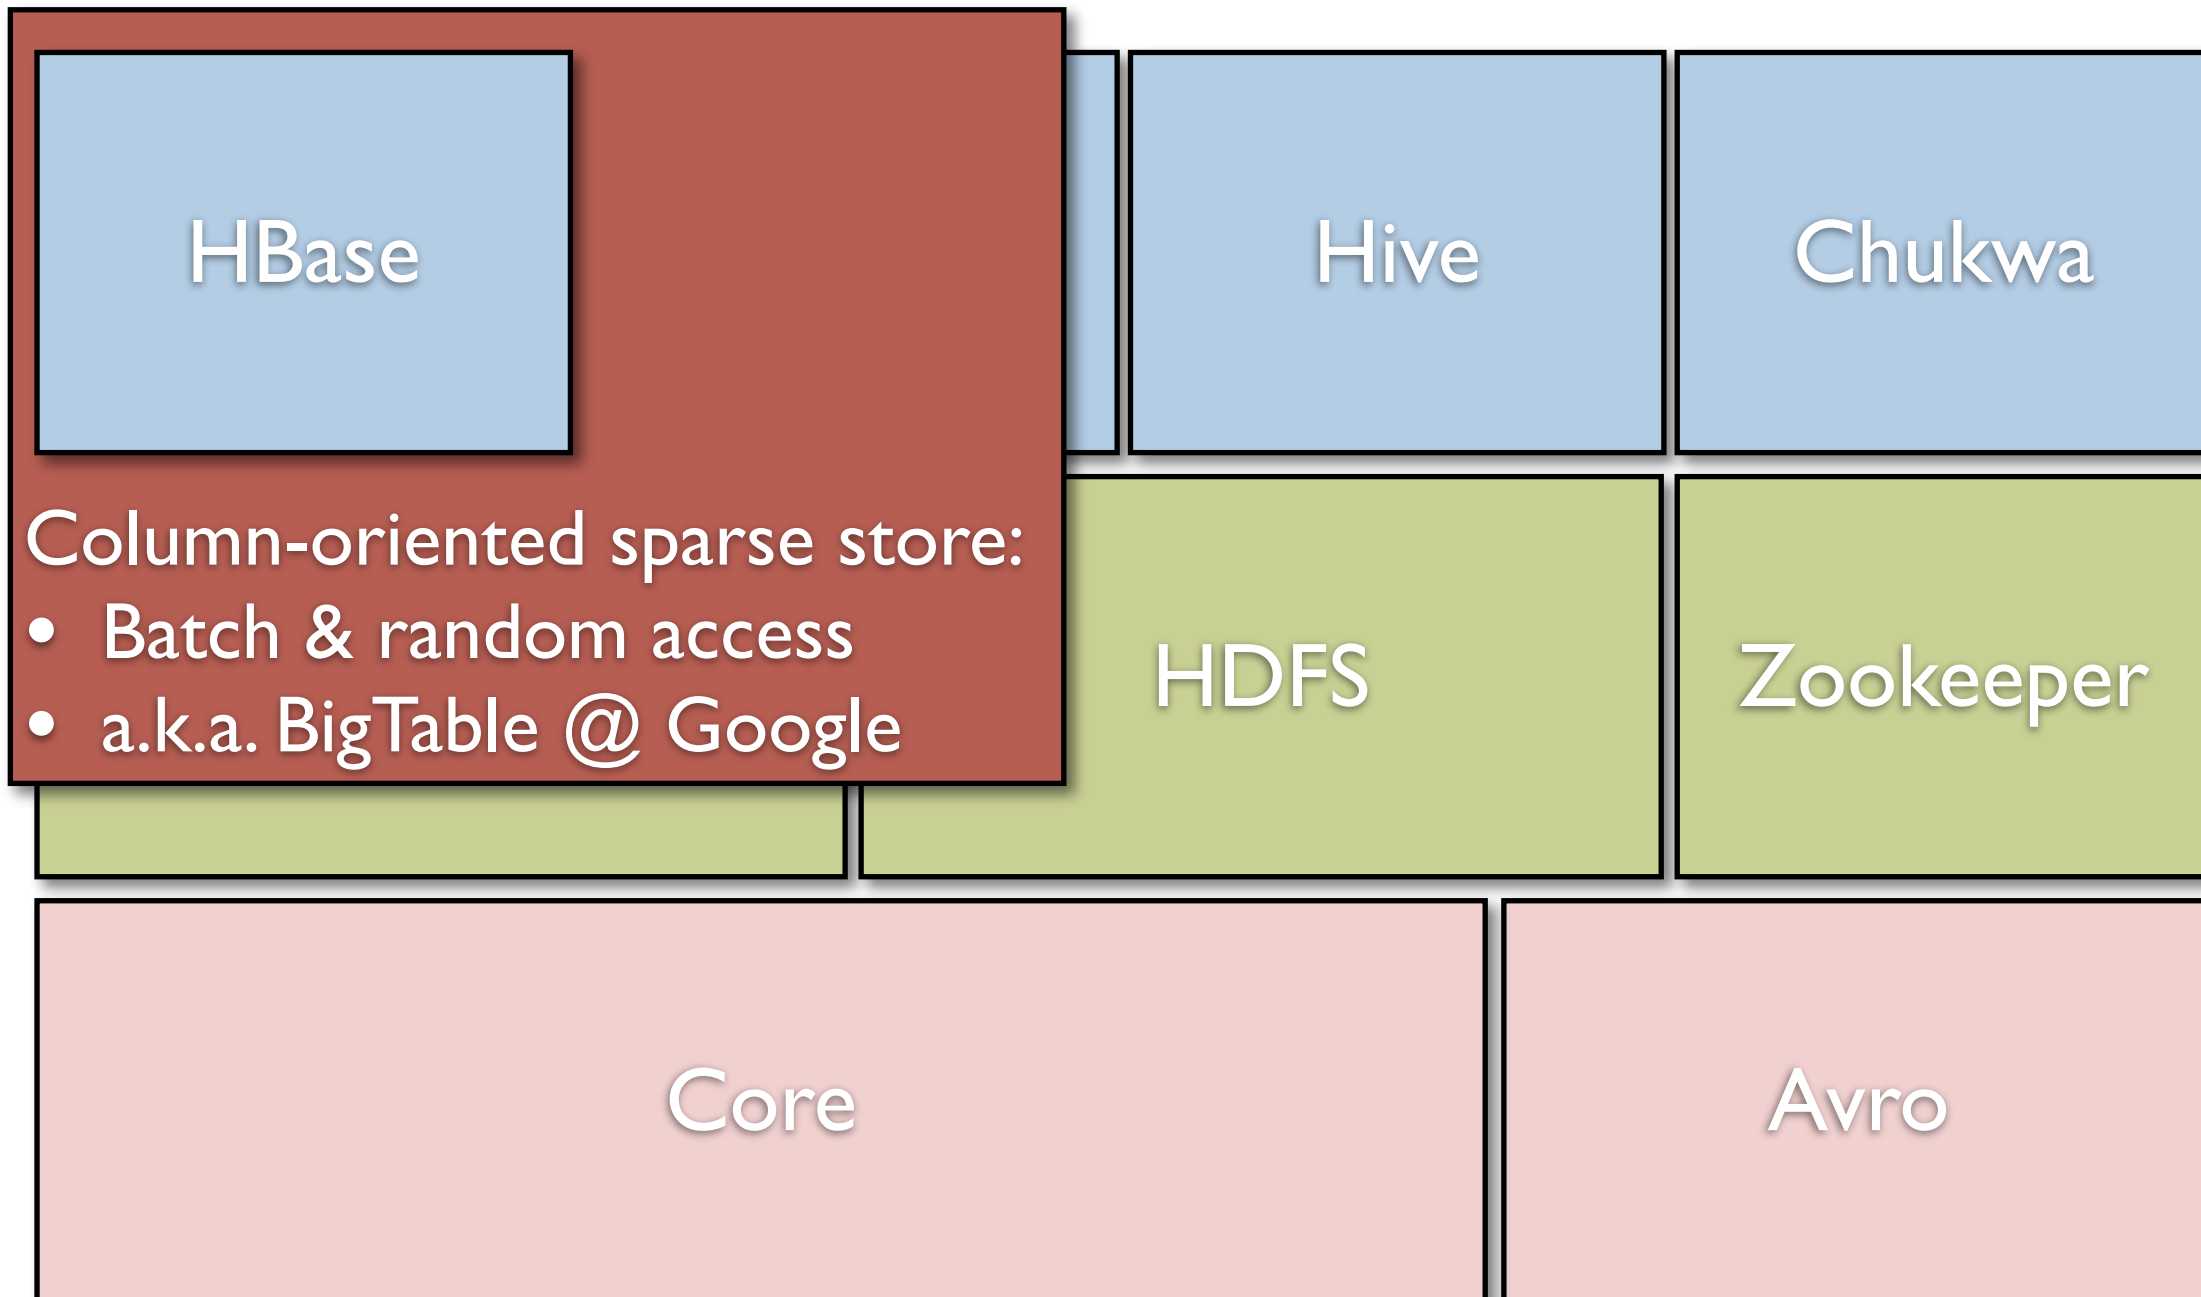
HADOOP



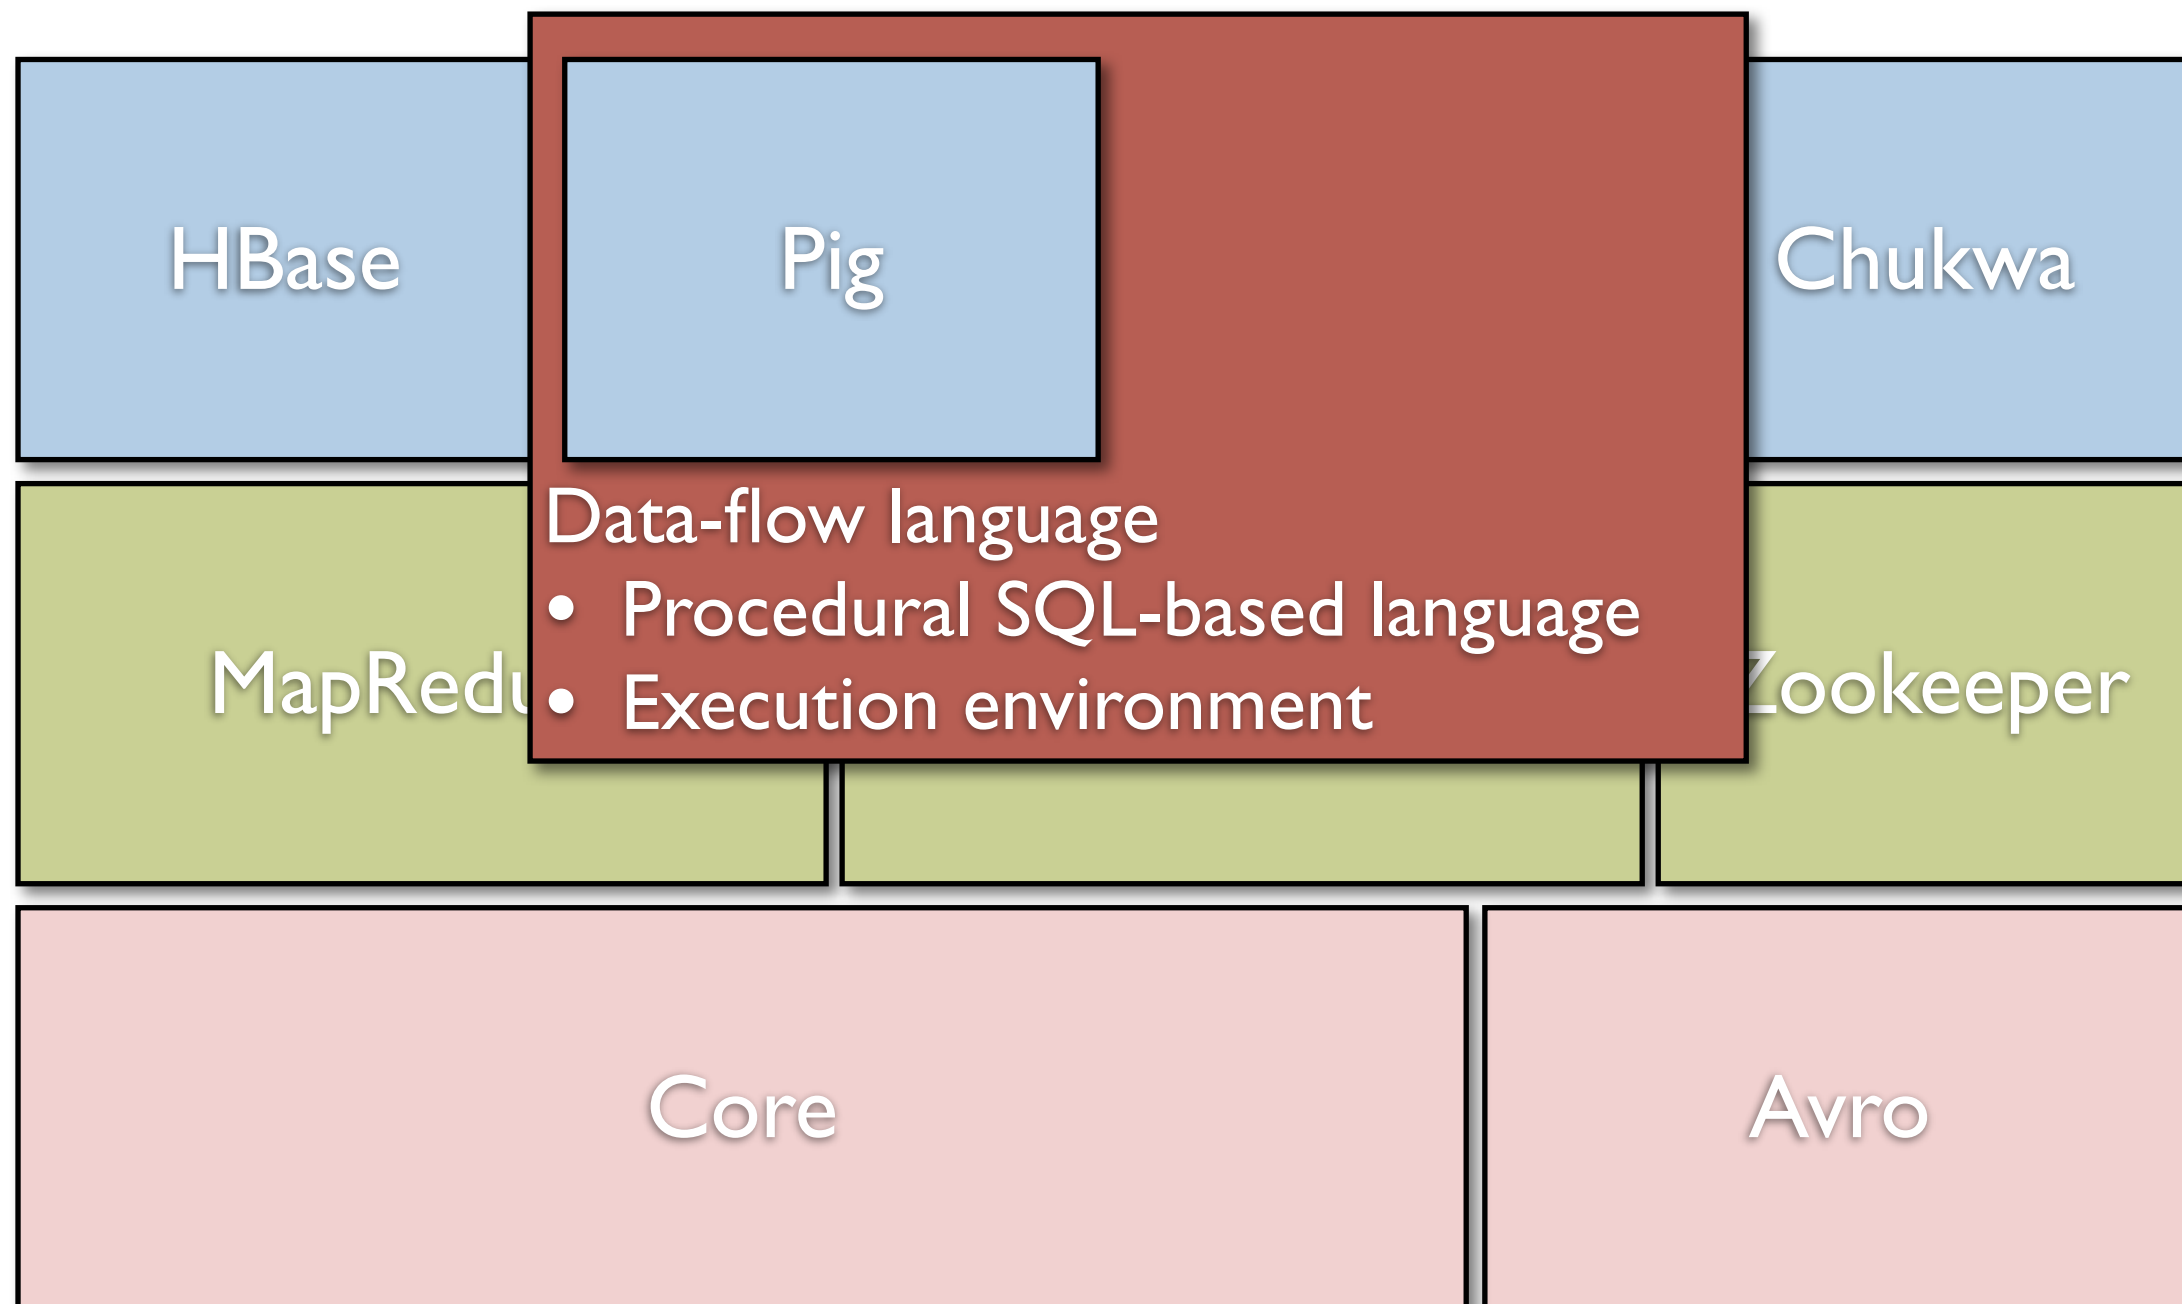
HADOOP



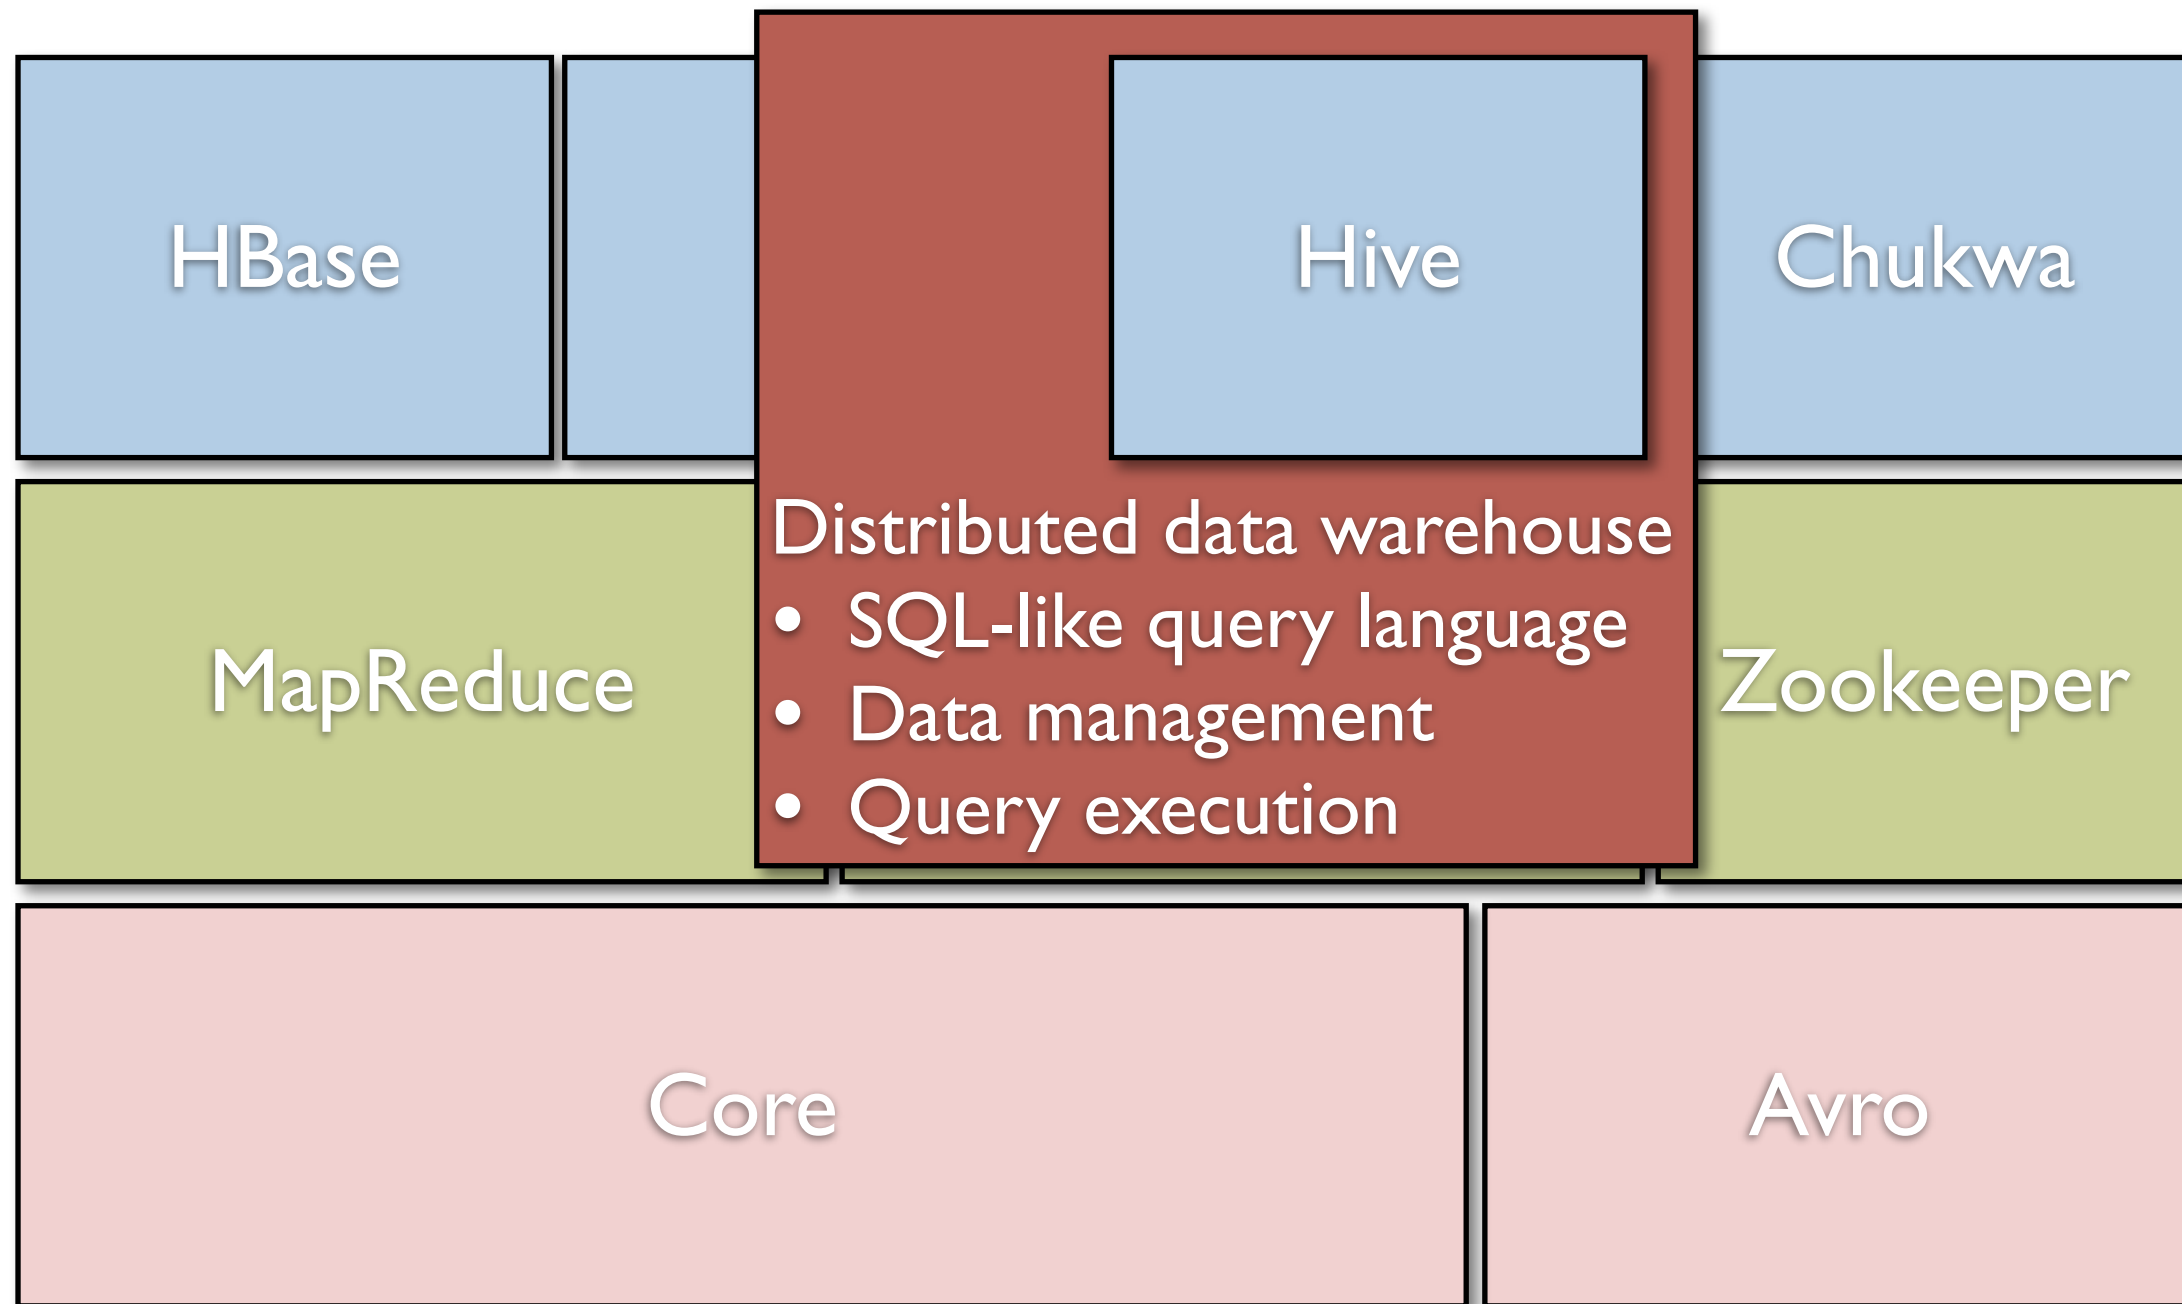




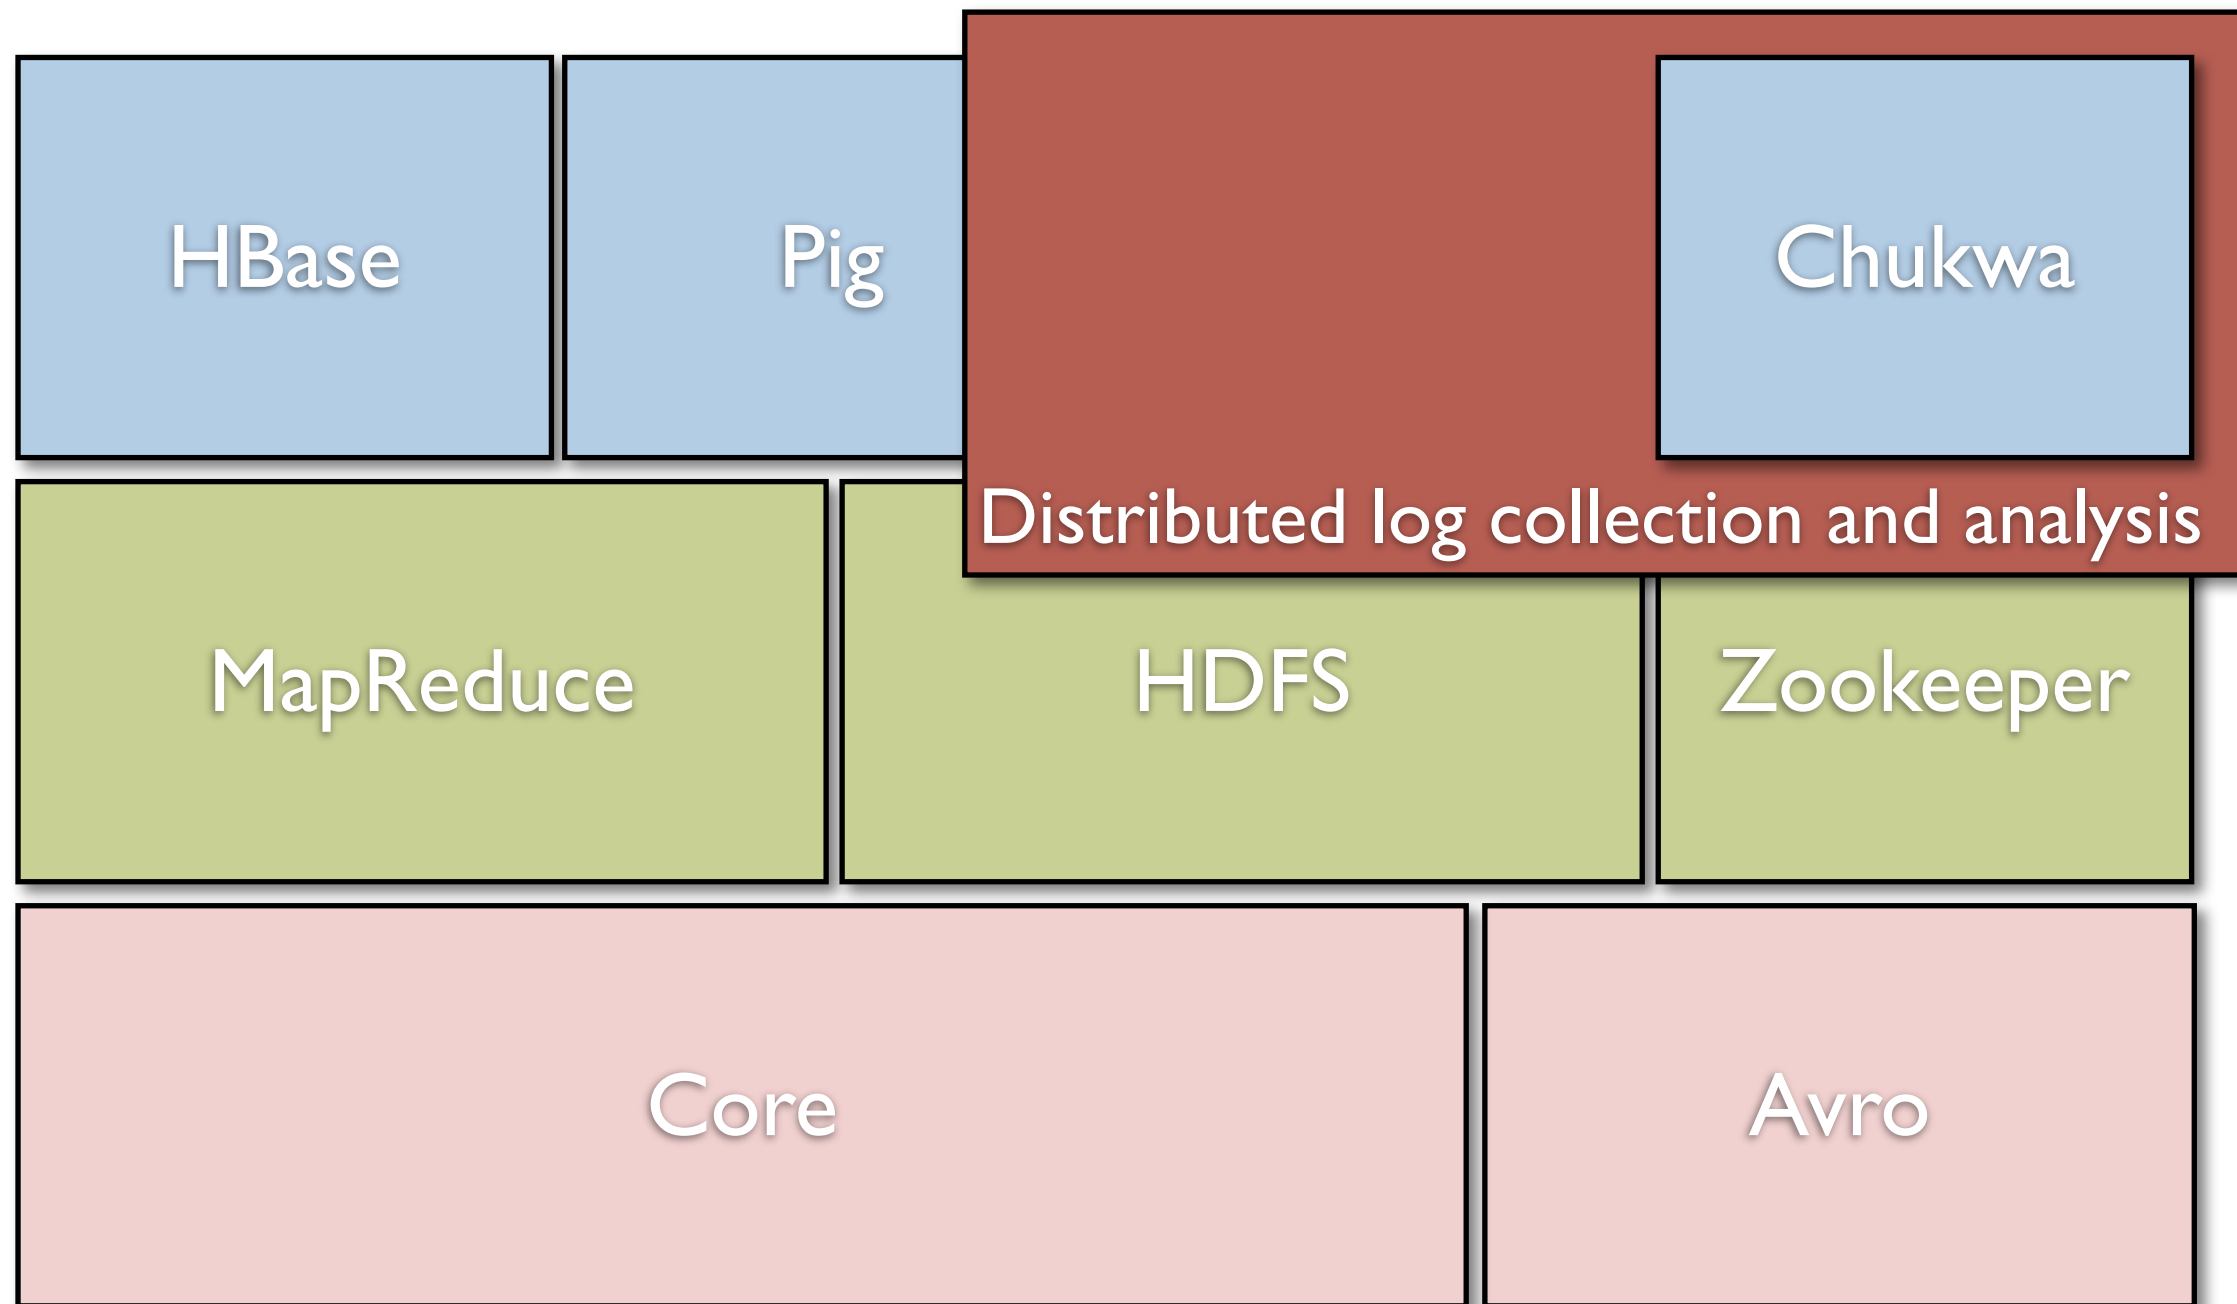
HADOOP



HADOOP




HADOOP



Basic HADOOP API (0.20.2)

- **Package org.apache.hadoop.mapreduce**
- **Class Mapper<KEYIN, VALUEIN, KEYOUT, VALUEOUT>**
 - void setup(Mapper.Context context)
 - void cleanup(Mapper.Context context)
 - void map(KEYIN key, VALUEIN value, Mapper.Context context)
 - output is generated by invoking context.collect(key, value);
- **Class Reducer<KEYIN, VALUEIN, KEYOUT, VALUEOUT>**
 - void setup(Reducer.Context context)
 - void cleanup(Reducer.Context context)
 - void reduce(KEYIN key, Iterable<VALUEIN> values, Reducer.Context context)
 - output is generated by invoking context.collect(key, value);
- **Class Partitioner<KEY, VALUE>**
 - abstract int getPartition(KEY key, VALUE value, int numPartitions)

- Package `org.apache.hadoop.io`



Defines a de/serialization protocol

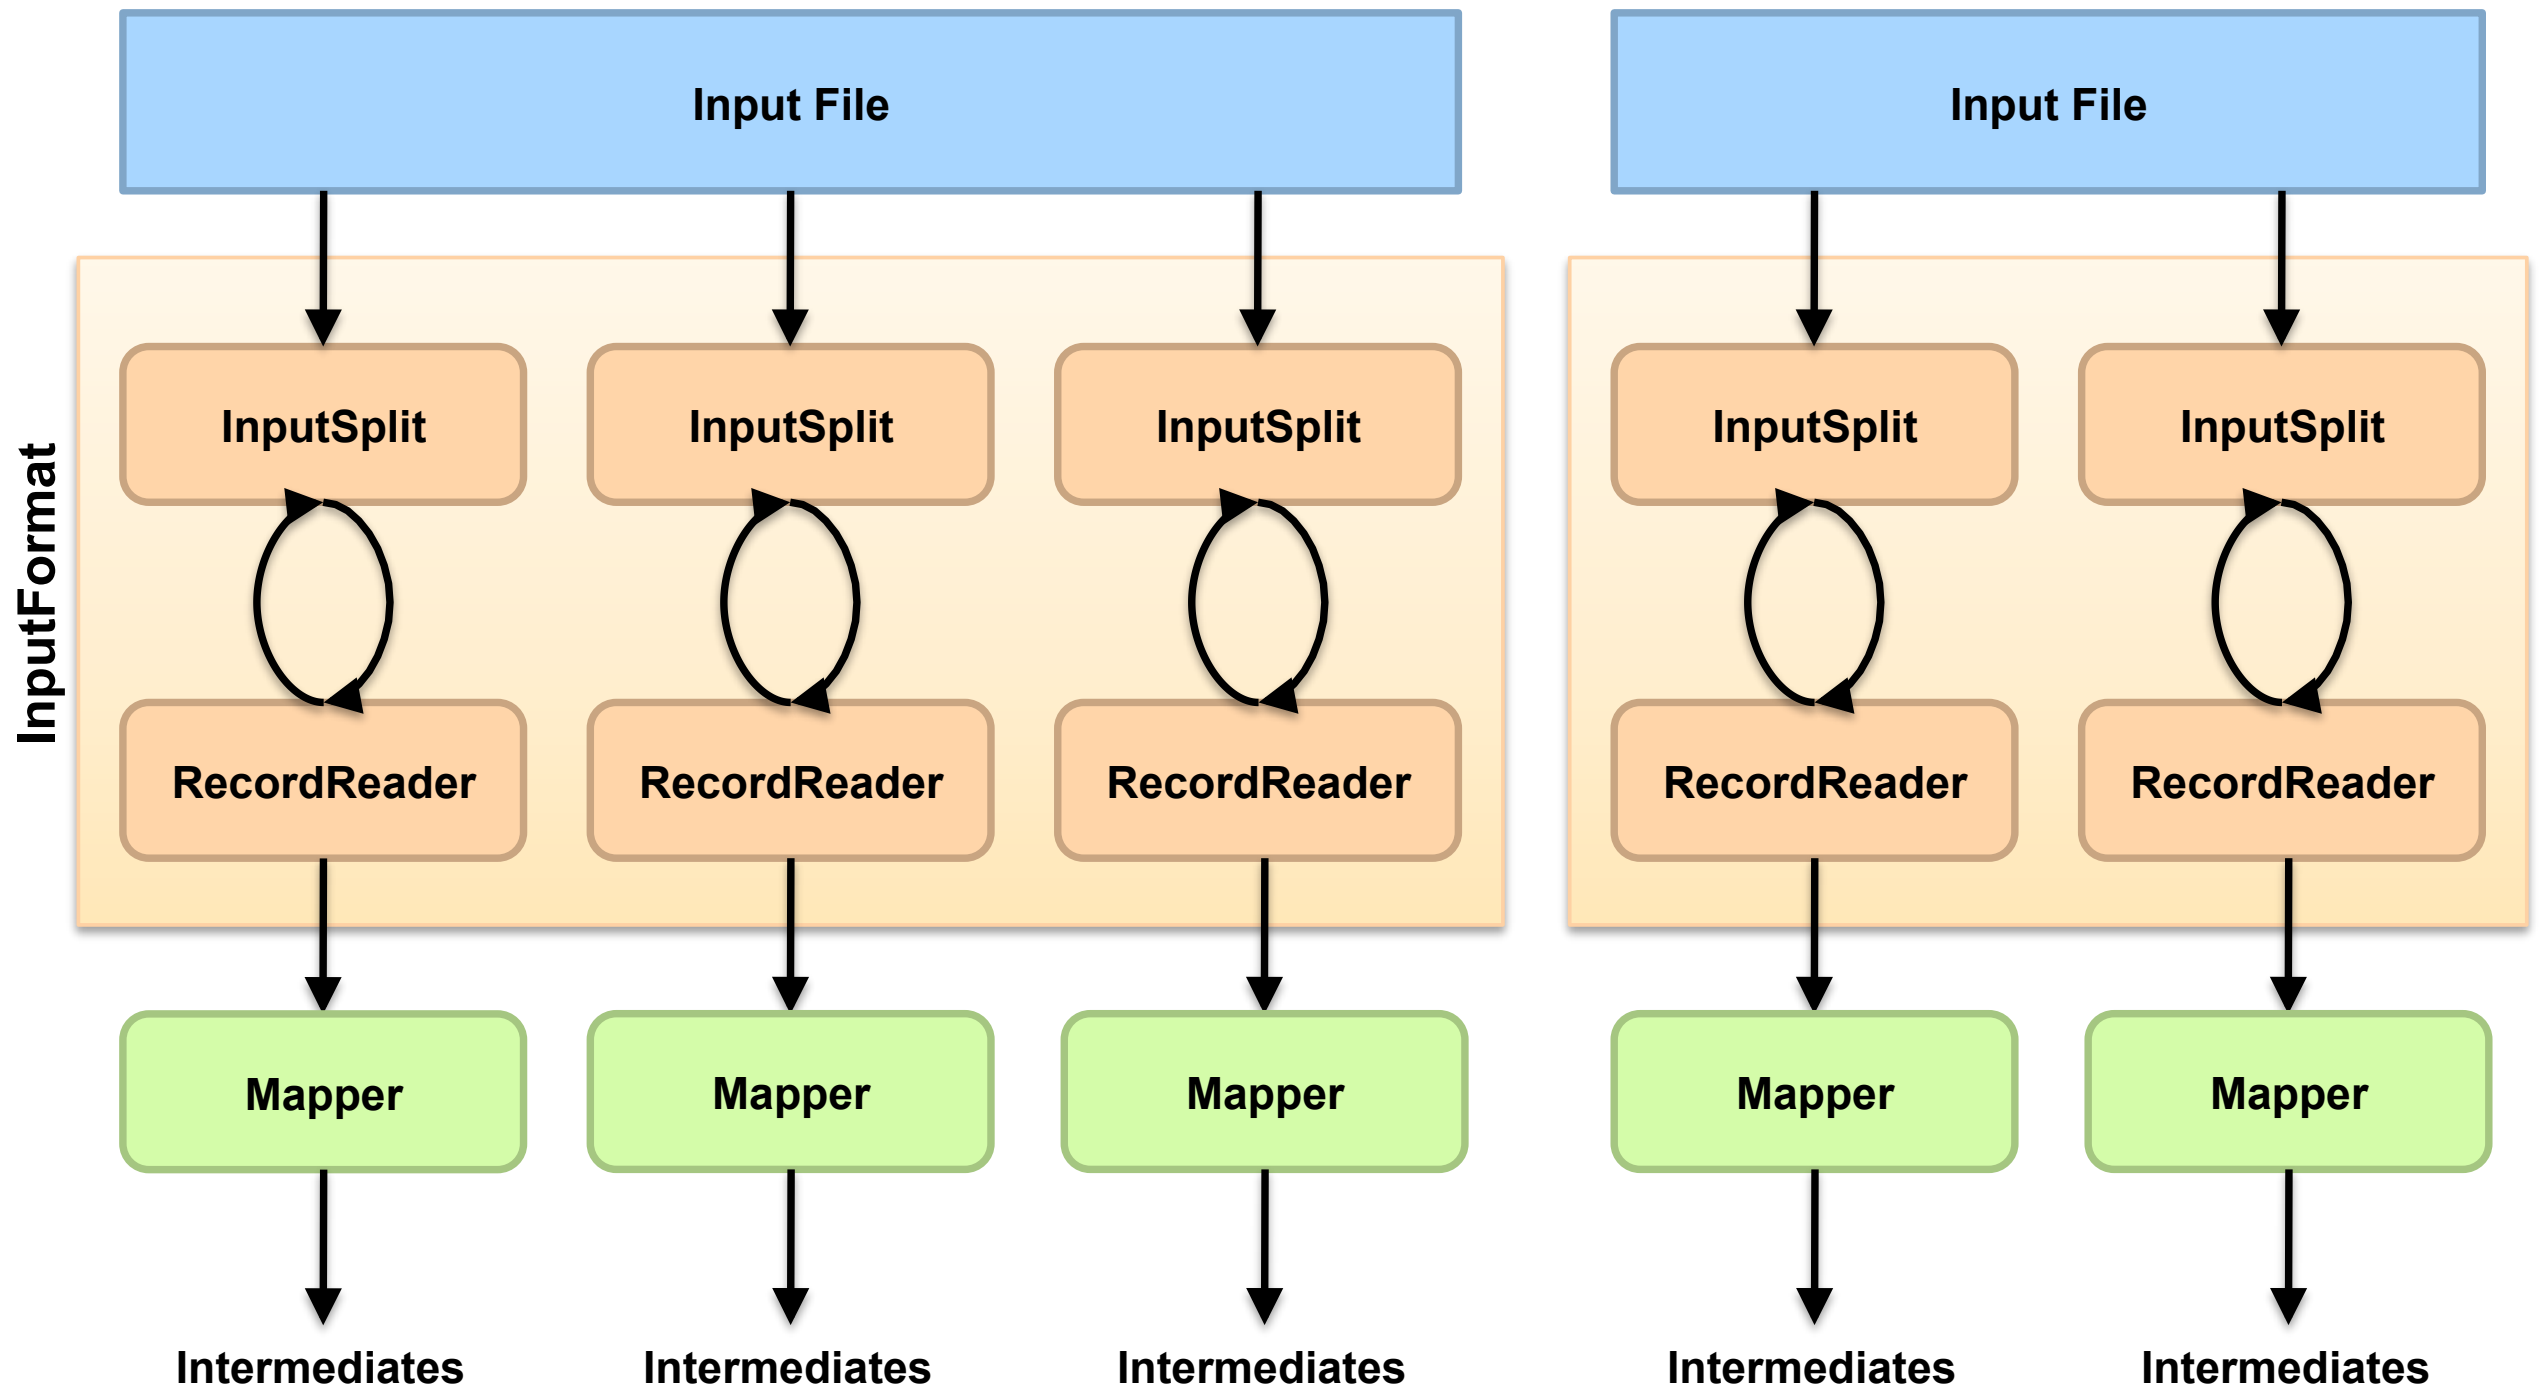
Any key or value type in the Hadoop Map-Reduce framework implements this interface

WritableComparables can be compared to each other, typically via Comparators

Any type which is to be used as a key in the Hadoop Map-Reduce framework should implement this interface

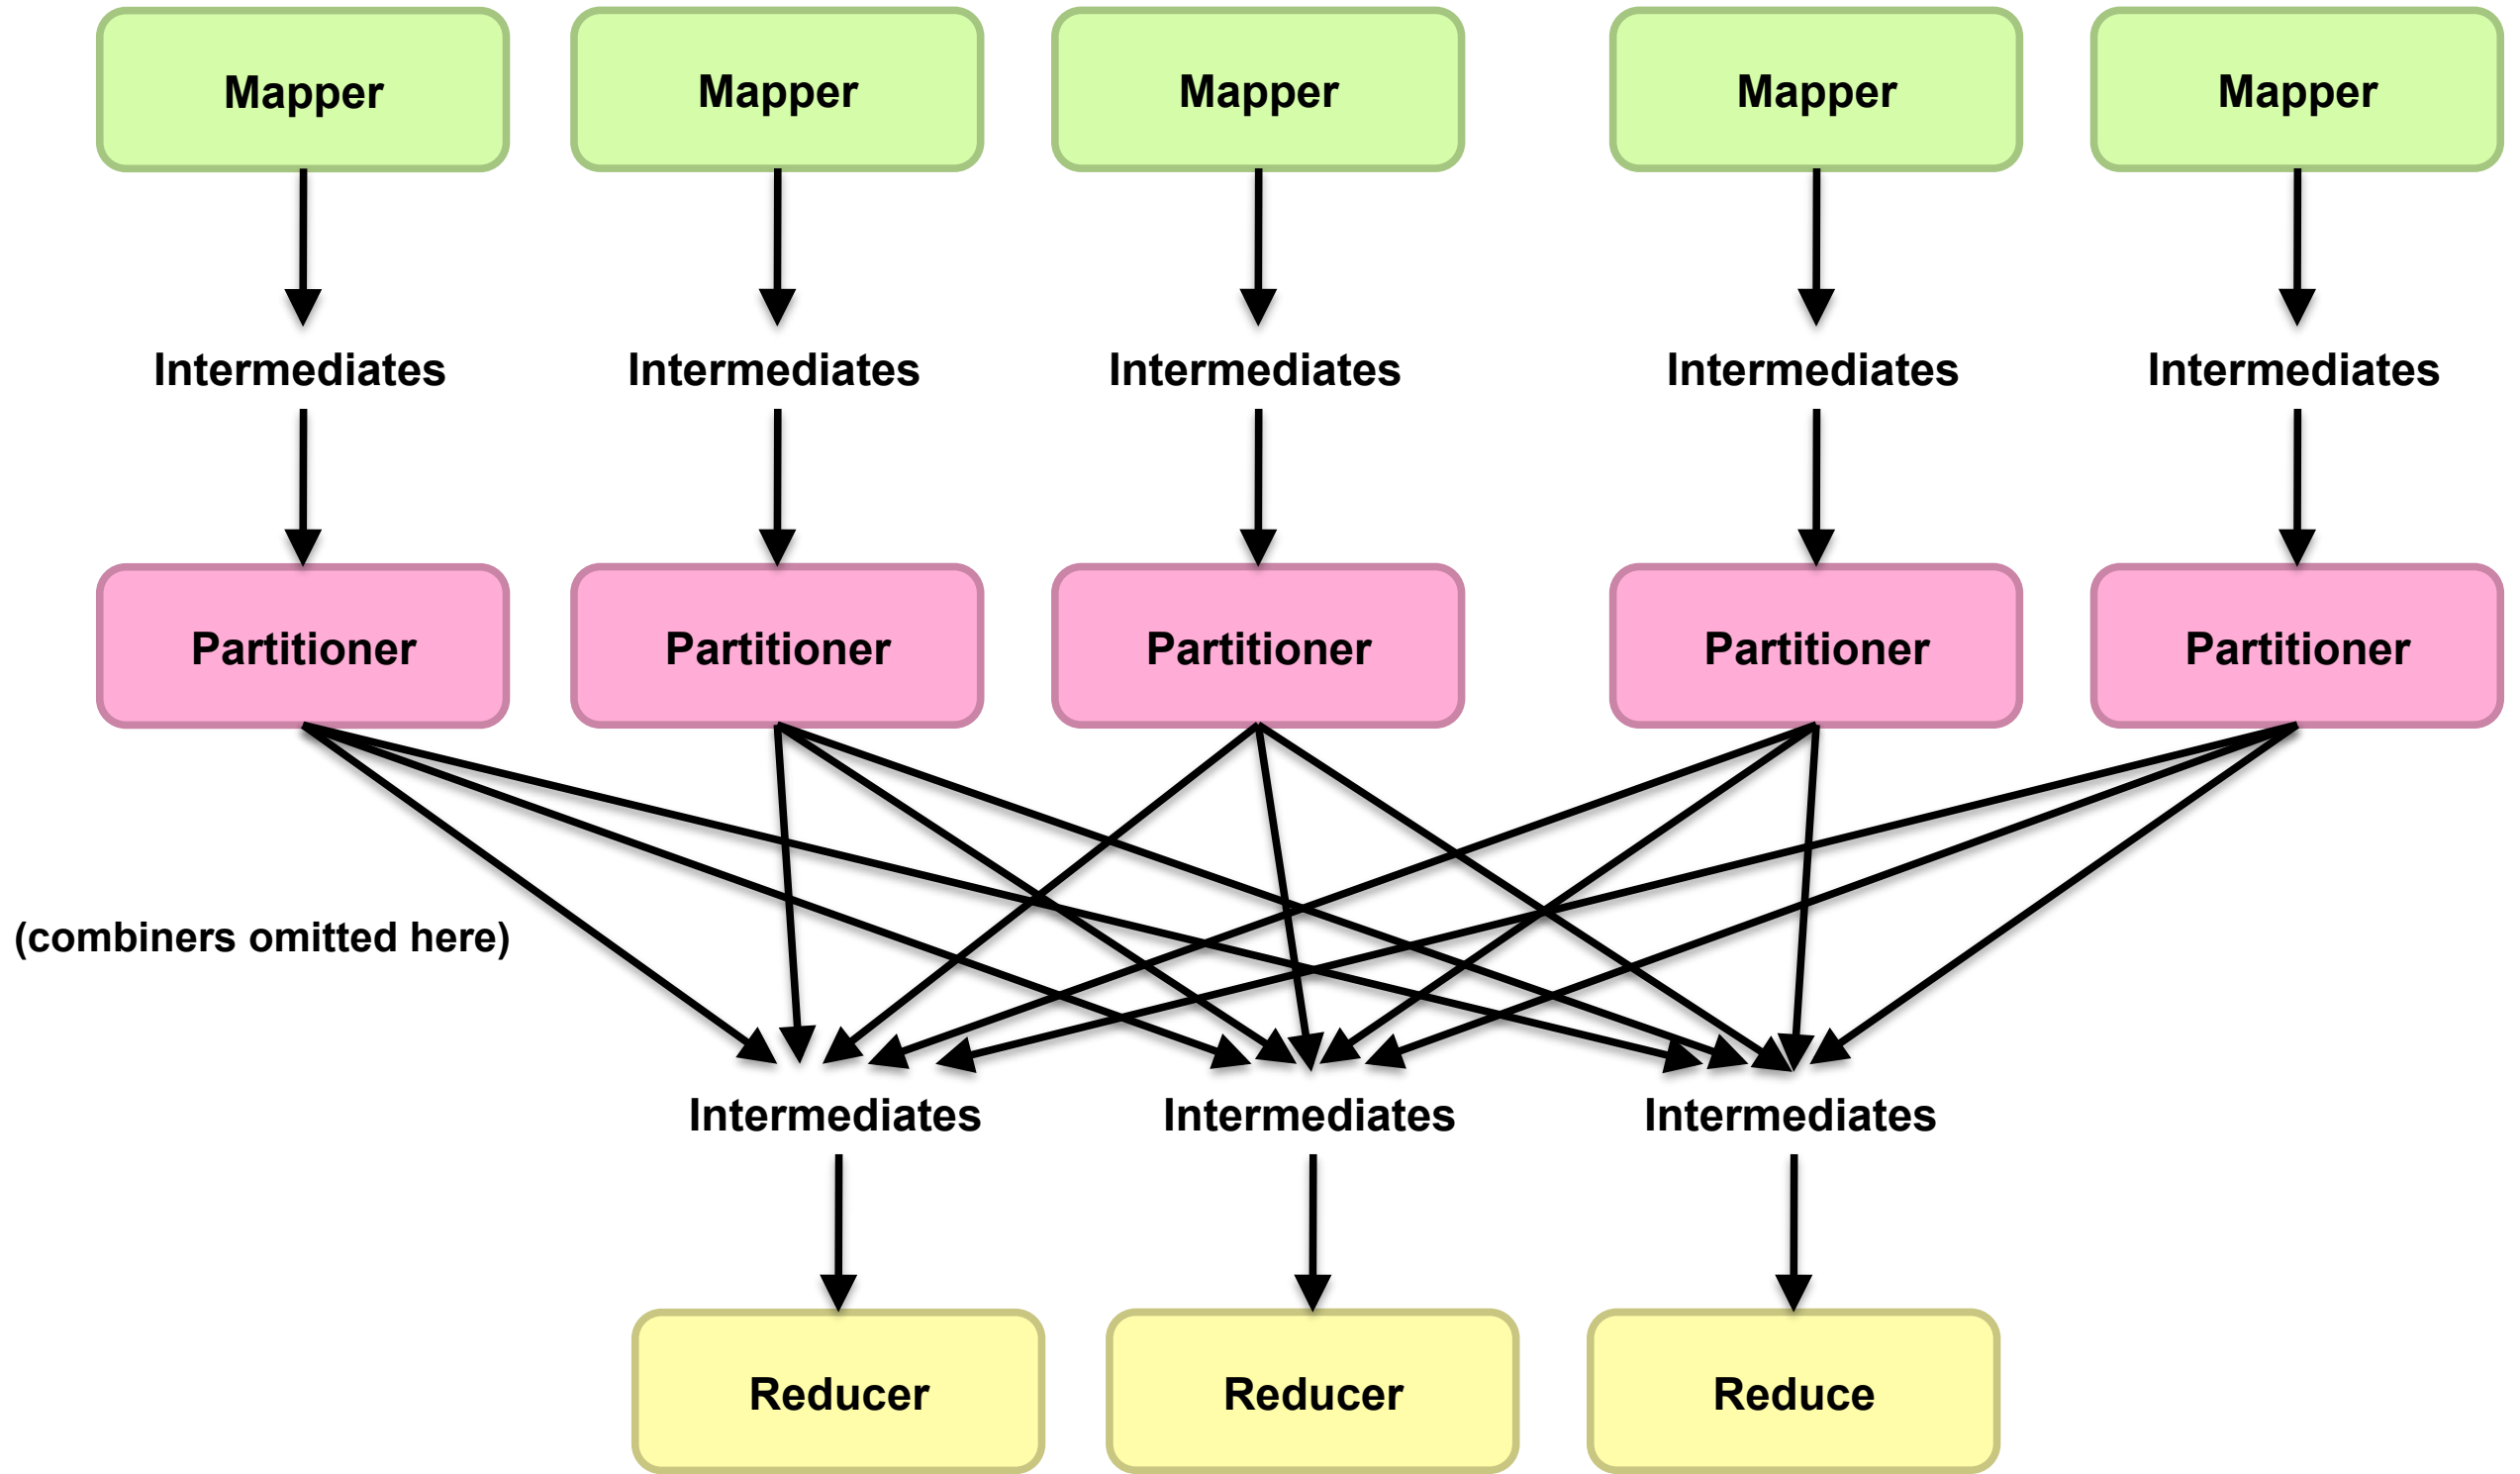
Concrete classes for common data types

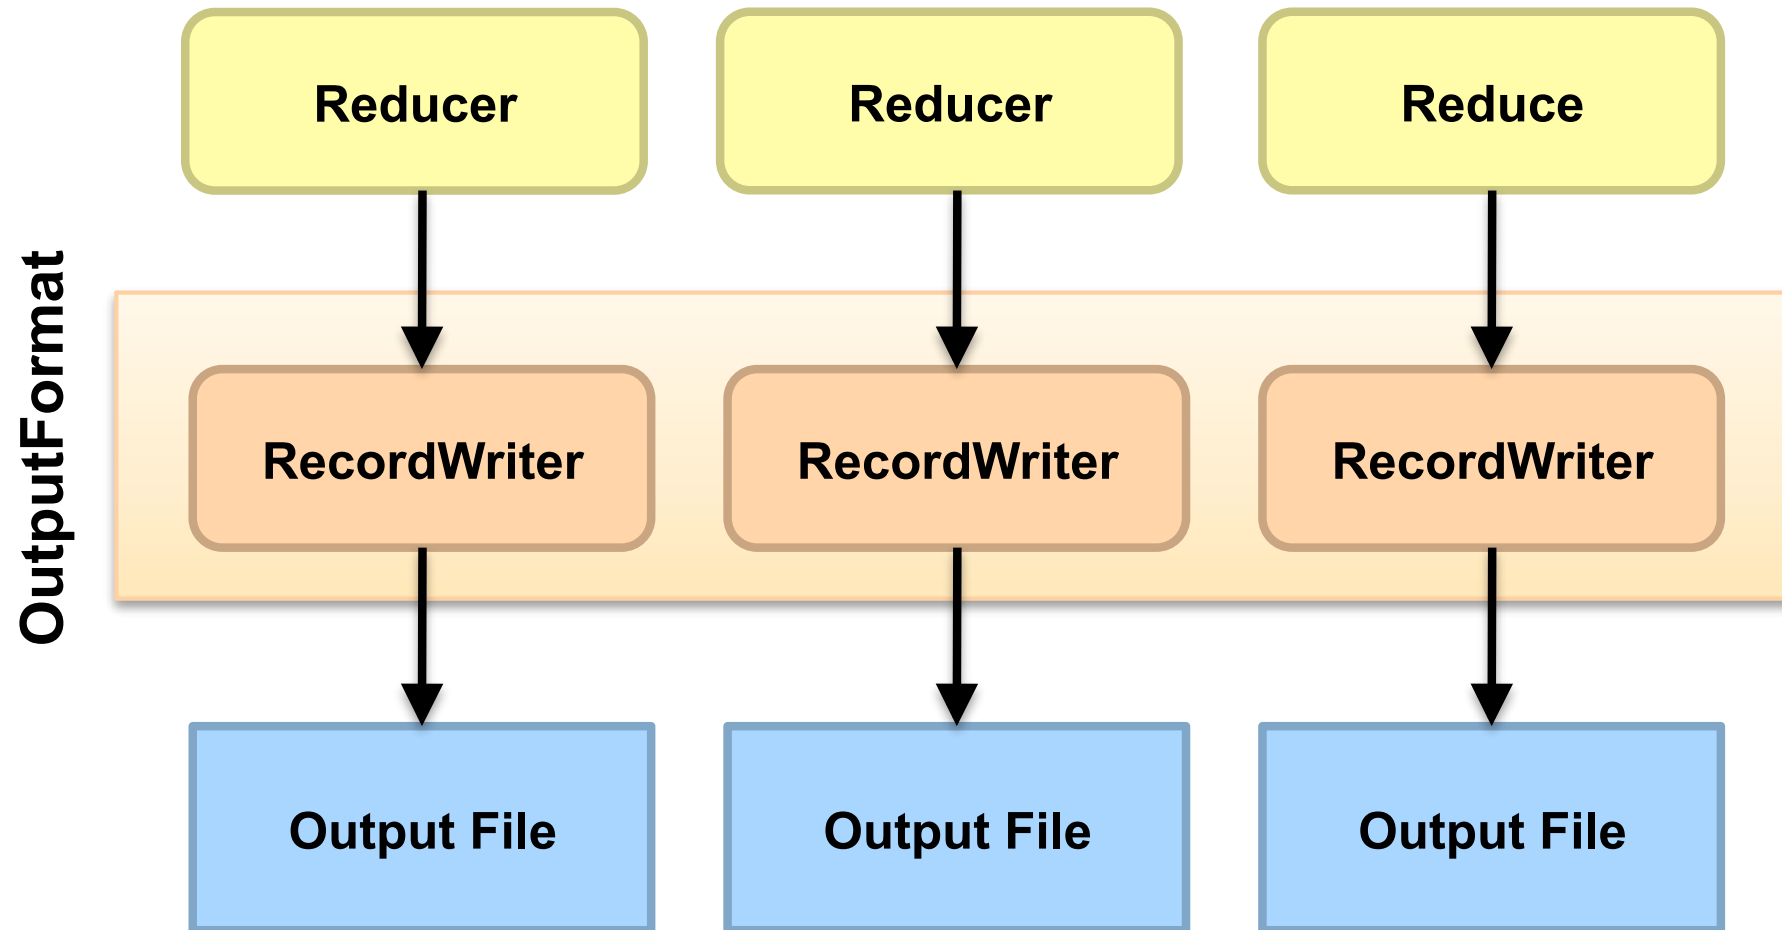
Hadoop Dataflow (I)



- Data sets are specified by **InputFormats**
 - Defines input data (e.g., a directory)
 - Identifies partitions of the data that form an **InputSplit**, each of which will be assigned to a mapper
 - Provide the **RecordReader** implementation to extract (k, v) records from the input source
- Base class implementation is **FileInputFormat**
 - Will read all files out of a specified directory and send them to the mappers
 - **TextInputFormat** – Treats each '\n'-terminated line of a file as a value
 - **KeyValueTextInputFormat** – Maps '\n'- terminated text lines of “k SEP v”
 - **SequenceFileInputFormat** – Binary file of (k, v) pairs with some add'l metadata
 - **SequenceFileAsTextInputFormat** – Same, but maps (k.toString(), v.toString())

Hadoop Dataflow (II)





- Data sets are specified by **OutputFormats**
 - Analogous to InputFormat
- Base class implementation is **FileOutputFormat**
 - TextOutputFormat – Writes “key val\n” strings to output file
 - SequenceFileOutputFormat – Uses a binary format to pack (k, v) pairs
- Other implementation is **NullOutputFormat**
 - Discards output to /dev/null

Hadoop Shuffle & Sort

- **Map Side**

- Mapper outputs are buffered in memory in a circular buffer
- When buffer reaches threshold, contents are “spilled” to disk
- Spills are merged in a single partitioned file (sorted within each partition)
- Combiners run here

- **Reduce Side**

- Firstly, mapper outputs are copied over to the reducer machine
- “Sort” is a multi-pass merge of map outputs (in memory and on disk)
- Combiners run here
- Final merge pass goes directly into reducer

- **Probably the most complex aspect of the framework!**