

# MapReduce Online

*Tyson Condie  
Neil Conway  
Peter Alvaro  
Joseph M. Hellerstein  
Khaled Elmeleegy  
Russell Sears*

Electrical Engineering and Computer Sciences  
University of California at Berkeley

Technical Report No. UCB/EECS-2009-136

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-136.html>

October 9, 2009



Copyright © 2009, by the author(s).  
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

#### Acknowledgement

We would like to thank Kuang Chen, Akshay Krishnamurthy, and Mosharaf Chowdhury for their helpful comments. This research is supported by grants from the National Science Foundation and the Natural Sciences and Engineering Research Council of Canada.

# MapReduce Online

Tyson Condie, Neil Conway, Peter Alvaro, Joseph M. Hellerstein  
*UC Berkeley*

Khaled Elmeleegy, Russell Sears  
*Yahoo! Research*

## Abstract

MapReduce is a popular framework for data-intensive distributed computing of batch jobs. To simplify fault tolerance, the output of each MapReduce task and job is *materialized* to disk before it is consumed. In this paper, we propose a modified MapReduce architecture that allows data to be *pipelined* between operators. This extends the MapReduce programming model beyond batch processing, and can reduce completion times and improve system utilization for batch jobs as well. We present a modified version of the Hadoop MapReduce framework that supports *online aggregation*, which allows users to see “early returns” from a job as it is being computed. Our Hadoop Online Prototype (*HOP*) also supports *continuous queries*, which enable MapReduce programs to be written for applications such as event monitoring and stream processing. *HOP* retains the fault tolerance properties of Hadoop, and can run unmodified user-defined MapReduce programs.

## 1 Introduction

MapReduce has emerged as a popular way to harness the power of large clusters of computers. MapReduce allows programmers to think in a *data-centric* fashion: they focus on applying transformations to sets of data records, and allow the details of distributed execution, network communication, coordination and fault tolerance to be handled by the MapReduce framework.

The MapReduce model is typically applied to large batch-oriented computations that are concerned primarily with time to job completion. The Google MapReduce framework [6] and open-source Hadoop system reinforce this usage model through a batch-processing implementation strategy: the entire output of each map and reduce stage is *materialized* to stable storage before it can be consumed by the next stage, or produce output. Batch materialization allows for a simple and elegant

checkpoint/restart fault tolerance mechanism that is critical in large deployments, which have a high probability of slowdowns or failures at worker nodes.

We propose a modified MapReduce architecture in which intermediate data is *pipelined* between operators, while preserving the programming interfaces and fault tolerance models of previous MapReduce frameworks. To validate this design, we developed the Hadoop Online Prototype (*HOP*), a pipelining version of Hadoop.

Pipelining provides several important advantages to a MapReduce framework, but also raises new design challenges. We highlight the potential advantages first:

- A downstream dataflow element can begin consuming data before a producer element has finished execution, which can increase opportunities for parallelism, improve utilization, and reduce response time. In Section 3.5, we show that *HOP* can achieve significantly lower job completion times than unmodified Hadoop.
- Since reducers begin processing data as soon as it is produced by mappers, they can generate and refine an approximation of their final answer during the course of execution. This technique, known as *online aggregation* [11], can reduce the turnaround time for data analysis by several orders of magnitude. We describe how we adapted online aggregation to our pipelined MapReduce architecture in Section 4.
- Pipelining widens the domain of problems to which MapReduce can be applied. In Section 5, we show how *HOP* can be used to support *continuous queries*: MapReduce jobs that run continuously, accepting new data as it arrives and analyzing it immediately. This allows MapReduce to be applied to domains such as system monitoring and stream processing.

Pipelining raises several design challenges. First, Google’s attractively simple MapReduce fault tolerance mechanism is predicated on the materialization of intermediate state. In Section 3.3, we show that this can coexist with pipelining, by allowing producers to periodically ship data to consumers in parallel with their materialization. A second challenge arises from the greedy communication implicit in pipelines, which is at odds with batch-oriented optimizations supported by “combiners”: map-side code that reduces network utilization by performing compression and pre-aggregation before communication. We discuss how the HOP design addresses this issue in Section 3.1. Finally, pipelining requires that producers and consumers are co-scheduled intelligently; we discuss our initial work on this issue in Section 3.4.

## 1.1 Structure of the Paper

In order to ground our discussion, we present an overview of the Hadoop MapReduce architecture in Section 2. We then develop the design of HOP’s pipelining scheme in Section 3, keeping the focus on traditional batch processing tasks, and demonstrating performance gains that pipelining can provide in that setting. Having explained HOP’s execution model, in Section 4 we show how it can support online aggregation for long-running jobs, and illustrate the potential benefits of that interface for MapReduce tasks. In Section 5 we describe our support for continuous MapReduce jobs over data streams, and demonstrate an example of near-real-time parallel system monitoring. Related and future work are covered in Sections 6 and 7.

## 2 Background

MapReduce is a programming model for performing transformations on large data sets [6]. In this section, we review the MapReduce programming model, and describe the salient features of Hadoop, a popular open-source implementation of MapReduce.

### 2.1 Programming Model

To use MapReduce, the programmer expresses their desired computation as a series of *jobs*. The input to a job is a list of *records* (key-value pairs). Each job consists of two steps: first, a user-defined *map* function is applied to each record to produce a list of intermediate key-value pairs. Second, a user-defined *reduce* function is called once for each distinct key in the map output, and passed the list of intermediate values associated with that key. The MapReduce framework automatically parallelizes the execution of these functions, and ensures fault tolerance.

```
public interface Mapper<K1, V1, K2, V2> {
    void map(K1 key, V1 value,
            OutputCollector<K2, V2> output);

    void close();
}
```

Figure 1: Map function interface.

Optionally, the user can supply a *combiner* function [6]. Combiners are similar to reduce functions, except that they are not passed *all* the values for a given key: instead, a combiner emits an output value that summarizes the input values it was passed. Combiners are typically used to perform map-side “preaggregation,” which reduces the amount of network traffic required between the map and reduce steps.

### 2.2 Hadoop Architecture

Hadoop is composed of *Hadoop MapReduce*, an implementation of MapReduce designed for large clusters, and the *Hadoop Distributed File System* (HDFS), a file system optimized for batch-oriented workloads such as MapReduce. In most Hadoop jobs, HDFS is used to store both the input to the map step and the output of the reduce step. Note that HDFS is *not* used to store intermediate results (e.g. the output of the map step): these are kept on each node’s local file system.

A Hadoop installation consists of a single master node and many worker nodes. The master, called the *JobTracker*, is responsible for accepting jobs from clients, dividing those jobs into *tasks*, and assigning those tasks to be executed by worker nodes. Each worker runs a *TaskTracker* process that manages the execution of the tasks currently assigned to that node. Each TaskTracker has a fixed number of slots for executing tasks (two maps and two reduces by default). A heartbeat protocol between each TaskTracker and the JobTracker is used to update the JobTracker’s bookkeeping of the state of running tasks, and drive the scheduling of new tasks: if the JobTracker identifies free TaskTracker slots, it will schedule further tasks on the TaskTracker.

### 2.3 Map Task Execution

Each map task is assigned a portion of the input file called a *split*. By default, a split contains a single HDFS block (64MB by default), so the size of the input file determines the number of map tasks.

The execution of a map task is divided into two phases.

1. The *map* phase reads the task’s split from HDFS, parses it into records (key/value pairs), and applies

the map function to each record.

2. After the map function has been applied to each input record, the *commit* phase registers the final output with the TaskTracker, which then informs the JobTracker that the task has finished executing.

Figure 1 contains the interface that must be implemented by user-defined map functions. After the *map* function has been applied to each record in the split, the *close* method is invoked.

The third argument to the *map* method specifies an *OutputCollector* instance, which accumulates the output records produced by the map function. The output of the map step is consumed by a reduce function, so the *OutputCollector* must store the map output in a format that is easy for the reducer to consume. Intermediate keys are assigned to reducers by applying a partitioning function, so the *OutputCollector* applies that function to each key produced by the map function, and stores each record and partition number in an in-memory buffer. The *OutputCollector* is responsible for spilling this buffer to disk when it reaches capacity.

A spill of the in-memory buffer involves first sorting the records in the buffer by partition number and then by key. The buffer content is written to the local file system as an index file and a data file (Figure 2). The index file points to the offset of each partition in the data file. The data file contains only the records, which are sorted by the key within each partition segment.

After a map task has applied the map function to each input record, it enters the *commit* phase. To generate the task's final output, the in-memory buffer is flushed to disk, and all of the spill files generated during the map phase are merged into a single data and index file. These output files are registered with the TaskTracker before the task completes. The TaskTracker will read these files when servicing requests from reduce tasks.

## 2.4 Reduce Task Execution

The execution of a reduce task is divided into three phases.

1. The *shuffle* phase fetches the reduce task's input data. Each reduce task is assigned a partition of the key range produced by the map step, so the reduce task must fetch the content of this partition from every map task's output.
2. The *sort* phase groups records with the same key together.
3. The *reduce* phase applies the user-defined reduce function to each key and corresponding list of values.

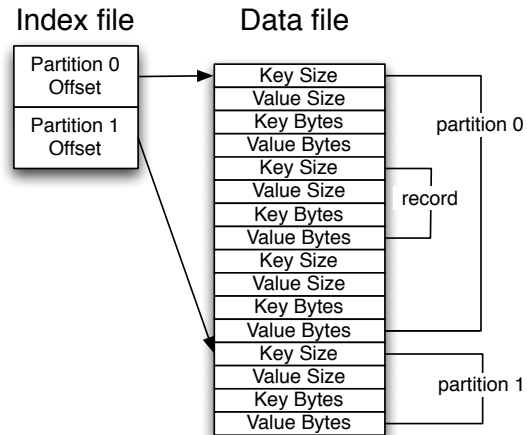


Figure 2: Map task index and data file format (2 partition/reduce case).

```
public interface Reducer<K2, V2, K3, V3> {
    void reduce(K2 key, Iterator<V2> values,
               OutputCollector<K3, V3> output);

    void close();
}
```

Figure 3: Reduce function interface.

In the *shuffle* phase, a reduce task fetches data from each map task by issuing HTTP requests to a configurable number of TaskTrackers at once (5 by default). The JobTracker relays the location of every TaskTracker that hosts map output to every TaskTracker that is executing a reduce task. In traditional batch-oriented Hadoop, a reduce task cannot fetch the output of a map task until the map has finished executing and committed its final output to disk.

After receiving its partition from all map outputs, the reduce task enters the *sort* phase. The map output for each partition is already sorted by key. The reduce task merges these runs together to produce a single run that is sorted by the key. The task then enters the *reduce* phase, in which it invokes the user-defined reduce function for each distinct key in sorted order, passing it the associated list of values. The output of the reduce function is written to a temporary location on HDFS. After the reduce function has been applied to each key in the reduce task's partition, the task's HDFS output file is atomically renamed from its temporary location to its final location.

In this design, the output of both map and reduce tasks is written to disk before it can be consumed. This is particularly expensive for reduce tasks, because their output is written to HDFS. By default, this requires a synchronous write operation that must store three copies of

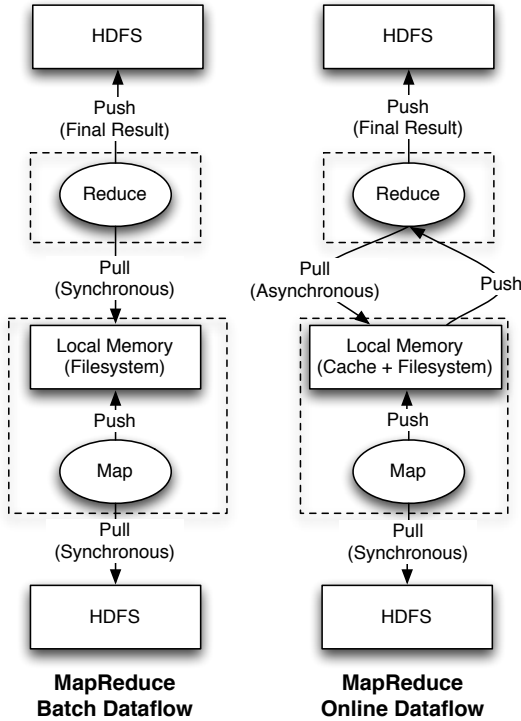


Figure 4: Hadoop dataflow for batch (left) and pipelined (right) processing of MapReduce computations.

each output block on different nodes (to ensure fault tolerance).

Output materialization simplifies fault tolerance, because it reduces the amount of state that must be restored to consistency after a node failure. If any task (either map or reduce) fails, the JobTracker simply schedules a new task to perform the same work as the failed task. Since a task never exports any data other than its final answer, no further recovery steps are needed.

### 3 Pipelined MapReduce

In this section we discuss our extensions to Hadoop to support pipelining. We focus here on batch-processing tasks, postponing the discussion of the changes required for online aggregation and continuous queries until the next sections.

Figure 4 depicts the dataflow of two MapReduce implementations. The dataflow on the left corresponds to the output materialization approach used by stock Hadoop; the dataflow on the right allows pipelining. In the remainder of this section, we present our design and implementation for the pipelined Hadoop dataflow. We describe how our design supports fault tolerance (Section 3.3), and discuss the interaction between pipelining and task scheduling (Section 3.4). We compare the

performance of stock Hadoop with our pipelining implementation in Section 3.5.

## 3.1 Pipelining Within A Job

As described in Section 2.4, reduce tasks traditionally issue HTTP requests to *pull* their output from each TaskTracker. This means that map task execution is completely decoupled from reduce task execution. To support pipelining, we modified the map task to instead *push* data to reducers as it is produced. To give an intuition for how this works, we begin by describing a straightforward pipelining design, and then discuss the changes we had to make to achieve good performance.

### 3.1.1 Naïve Pipelining

In our naïve implementation, we want to directly connect mappers to reducers and pipe data between them. When a client submits a new job to Hadoop, the JobTracker assigns the map and reduce tasks associated with the job to the available TaskTracker slots. For purposes of discussion, let us assume that there are enough free slots to assign all the tasks for each job. The JobTracker also communicates the location of each map task to every reduce task. We modified Hadoop so that each reduce task contacts every map task upon initiation of the job, and opens a socket which will be used to pipeline the output of the map function. As each map output record is produced, the mapper determines which partition (reduce task) the record should be sent to, and immediately sends it via the appropriate socket.

A reduce task accepts the pipelined data it receives from each map task and stores it in an in-memory buffer, spilling sorted runs of the buffer to disk as needed. Once the reduce task learns that every map task has completed, it performs a final merge of all its sorted runs and applies the user-defined reduce function as normal, writing the output to HDFS.

### 3.1.2 Refinements

While the algorithm described above is straightforward, it suffers from several practical problems. First, it is possible that there will not be enough slots available to schedule every task in a new job. Opening a socket between every map and reduce task also requires a large number of TCP connections. A simple tweak to the naïve design solves both problems: if a reduce task has not yet been scheduled, any map tasks that produce records for that partition simply write them to disk. Once the reduce task is assigned a slot, it can then fetch the records from the map task, as in stock Hadoop. To reduce the number of concurrent TCP connections, each reducer can be

configured to pipeline data from a bounded number of mappers at once; the reducer will pull data from the remaining map tasks in the traditional Hadoop manner.

Our initial pipelining implementation suffered from a second problem: the user-defined map function was invoked by the same thread that wrote output records to the pipeline sockets. This meant that if a network I/O blocked (e.g. because the reducer was over-utilized), the mapper was prevented from doing useful work. Pipeline stalls should not prevent a map task from making progress—especially since, once the task has completed, it frees a TaskTracker slot that can be used for other purposes. We solved this problem by running the map function in a separate thread that stores its output in an in-memory buffer, and then having another thread periodically send the contents of the buffer to the pipelining reducers.

### 3.1.3 Granularity of Map Output

Another problem with the naïve design is that it eagerly sends each record as soon as it is produced, which prevents the use of map-side combiners. Imagine a job where the reduce key has few distinct values (e.g., gender), and the reduce applies an aggregate function (e.g., count). As discussed in Section 2.1, combiners allow “map-side preaggregation”: by applying a reduce-like function to each distinct key at the mapper, network traffic can often be substantially reduced. By eagerly pipelining each record as it is produced, there is no opportunity for the map task to apply a combiner function.

A related problem is that eager pipelining moves some of the sorting work from the mapper to the reducer. Recall that in the blocking architecture, map tasks generate sorted spill files: all the reduce task must do is merge together the pre-sorted map output for each partition. In the naïve pipelining design, map tasks send output records in the order in which they are generated, so the reducer must perform a full external sort. Because the number of map tasks typically far exceeds the number of reduces [6], moving more work to the reducer increased response time in our experiments.

We addressed these issues by modifying the in-memory buffer design we described in Section 3.1.2. Instead of sending the buffer contents to reducers directly, we instead wait for the buffer to grow to a threshold size. The mapper then applies the combiner function, sorts the output by partition and reduce key, and writes the buffer to disk, using the spill file format described in Section 2.3. A second thread monitors the spill files, and sends them to the pipelined reducers. If the reducers are able to keep up with the map task and the network is not a bottleneck, a spill file will be sent to a reducer very quickly after it has been produced (in which case, the

spill file will likely remain in the mapper machine’s kernel buffer cache). However, if the reducer begins to trail the mapper, the number of unsent spill files will grow. In this case, the mapper periodically applies the combiner function to the spill files, merging multiple spill files together into a single larger file. This has the effect of *adaptively* moving load from the reducer to the mapper or vice versa, depending on which node is the current bottleneck.

The connection between pipelining and adaptive query processing techniques has been observed elsewhere (e.g. [2]). The adaptive scheme outlined above is relatively simple, but we believe that adapting to feedback along pipelines has the potential to significantly improve the utilization of MapReduce clusters.

## 3.2 Pipelining Between Jobs

Many practical computations cannot be expressed as a single MapReduce job, and the outputs of higher-level languages like Pig [19] typically involve multiple jobs. In the traditional Hadoop architecture, the output of each job is written to HDFS in the reduce step, and then immediately read back from HDFS by the map step of the next job. In fact, the JobTracker cannot even schedule a consumer job until the producer job has completed, because scheduling a map task requires knowing the HDFS block locations of the map’s input split.

In our modified version of Hadoop, the reduce tasks of one job can optionally pipeline their output directly to the map tasks of the next job, sidestepping the need for expensive fault-tolerant storage in HDFS for what amounts to a temporary file. Unfortunately, the computation of the reduce function from the previous job and the map function of the next job cannot be overlapped: the final result of the reduce step cannot be produced until all map tasks have completed, which prevents effective pipelining. However, in the next sections we describe how online aggregation and continuous query pipelines can publish “snapshot” outputs that can indeed pipeline between jobs.

## 3.3 Fault Tolerance

Our pipelined Hadoop implementation is robust to the failure of both map and reduce tasks. To recover from map task failures, we added some bookkeeping to the reduce task to record which map task produced each pipelined spill file. To simplify fault tolerance, the reducer treats the output of a pipelined map task as “tentative” until the JobTracker informs the reducer that the map task has committed successfully. The reducer can merge together spill files generated by the same uncommitted mapper, but won’t combine those spill files with

the output of other map tasks until it has been notified that the map task has committed. Thus, if a map task fails, each reduce task can ignore any tentative spill files produced by the failed map attempt. The JobTracker will take care of scheduling a new map task attempt as in stock Hadoop.

If a reduce task fails and a new copy of the task is started, the new reduce instance must be sent all the input data that was sent to the failed reduce attempt. If map tasks operated in a purely pipelined fashion and discarded their output after sending it to a reducer, this would be difficult. Therefore, map tasks retain their output data, and write a complete output file to disk before committing as in normal Hadoop. This allows the map’s output to be reproduced if any reduce tasks fail. For batch jobs, the key advantage of our architecture is that reducers are not blocked waiting for the complete output of the task to be written to disk.

Our technique for recovering from map task failure is straightforward, but places a minor limit on the reducer’s ability to merge spill files. To avoid this, we envision introducing a “checkpoint” concept: as a map task runs, it will periodically notify the JobTracker that it has reached offset  $x$  in its input split. The JobTracker will notify any connected reducers; map task output that was produced before offset  $x$  can then be merged by reducers with other map task output as normal. To avoid duplicate results, if the map task fails, the new map task attempt resumes reading its input at offset  $x$ . This technique also has the benefit of reducing the amount of redundant work done after a map task failure.

### 3.4 Task Scheduling

The Hadoop JobTracker had to be retrofitted to be aware of the pipelined execution of jobs. Intra-job pipelining was partially handled by the baseline version of Hadoop, which co-schedules map and reduce tasks that are part of the same job. However, the JobTracker was not aware of inter-job dependencies. In stock Hadoop, a client that needs to submit a series of jobs (that possibly consist of a larger query) must do so in order of their dataflow dependencies. That is, a job that consumes the output of one or more other jobs cannot be submitted until the producer jobs have executed to completion.

Hadoop exports an interface to client applications for submitting jobs. We extended this interface to accept a list (pipeline) of jobs, where each job in the list depends on the job before it. The client interface traverses this list annotating each job with the identifier of the job that it depends on. The JobTracker scheduler looks for this annotation and co-schedules jobs with their dependencies, giving slot preference to “upstream” jobs over the “downstream” jobs they feed. As we note in Sec-

tion 7, there are many interesting options for scheduling pipelines or even DAGs of such jobs that we plan to investigate in future.

### 3.5 Performance Evaluation

We conducted a series of performance experiments using a 60-node cluster on Amazon EC2. One node executed the Hadoop JobTracker and the HDFS NameNode, while the remaining 59 nodes served as slaves for running the TaskTrackers and HDFS DataNodes. All nodes executed on “high-CPU medium” EC2 instances with 1.7GB of memory and 2 virtual cores. Each virtual core is the equivalent of a 2007-era 2.5Ghz Intel Xeon processor.

We began by measuring the performance of a simple MapReduce job that does not use a combiner. Sorting is commonly used as a benchmark for basic MapReduce performance, because of the implicit sort done by the reduce phase. We sorted 5.5GB of article text extracted from Wikipedia; each word from the text was parsed as a separate record. Figure 5 describes sort performance on the EC2 cluster using an HDFS block size of 128MB (yielding 40 map tasks). We configured the system to use 59 reducers. In each graph, we give the CDF of map and reduce task completion. The left and right graphs depict blocking and pipelined performance, respectively.

Pipelining dominates blocking for this configuration, in part because it achieves better cluster utilization: the reduce tasks in the blocking job were idle for the first 192 seconds of the experiment, whereas for the pipelined job, reducers began doing useful work within 20 seconds. Note that in a highly-utilized cluster, increased pipeline parallelism would not necessarily lead to an improvement in total throughput. However, these results suggest that pipelining can substantially reduce the response time of an individual job, which can often be important (e.g. quickly executing high-priority jobs).

## 4 Online Aggregation

Although MapReduce was originally designed as a batch-oriented system, it is often used for interactive data analysis: a user submits a job to extract information from a data set, and then waits to view the results before proceeding with the next step in the data analysis process. This trend has accelerated with the development of high-level query languages that are executed as MapReduce jobs, such as Hive [25], Pig [19], and Sawzall [22].

Traditional MapReduce implementations provide a poor interface for interactive data analysis, because they do not emit any output until the job has been executed to completion. However, in many cases, an interactive user would prefer a “quick and dirty” approximation over a correct answer that takes much longer to compute. In the






Figure 5: CDF of map and reduce task completion times for a sort job on 5.5GB of text extracted from Wikipedia. The total job runtimes were 927 seconds for blocking, and 610 seconds for pipelining.

database literature, online aggregation has been proposed to address this problem [11], but the batch-oriented nature of traditional MapReduce implementations makes these techniques difficult to apply. In this section, we show how we extended our pipelined Hadoop implementation to support online aggregation within a single job (Section 4.1) and between multiple jobs (Section 4.2). We show that online aggregation has a minimal impact on job completion times, and can often yield an accurate approximate answer long before the job has finished executing.

## 4.1 Single-Job Online Aggregation

In our pipelined version of Hadoop, the data records produced by each map task are incrementally sent to reduce tasks, shortly after each record is generated. To produce the final output of the job, the reduce function cannot be invoked until all the output of every map task has been produced. We can support online aggregation by simply applying the reduce function to the data that a reduce task has received so far. We call the output of such an intermediate reduce operation a *snapshot*.

Users would like to know how accurate a snapshot is: that is, how closely a snapshot resembles the final output of the job. Accuracy estimation is a hard problem even for simple SQL queries [14], and particularly hard for jobs where the map and reduce functions are opaque user-defined code. Hence, we report job *progress*, not accuracy: we leave it to the user (or their MapReduce code) to correlate progress to a formal notion of accuracy. We give a simple progress metric below.

Snapshots are computed periodically, as new data arrives at each reducer. The user specifies how often snapshots should be computed, using the progress metric as the unit of measure. For example, a user can request that a snapshot be computed when 25%, 50%, and 75% of the input has been seen. The user may also specify whether

to include data from tentative (unfinished) map tasks. This option does not affect the fault tolerance design described in Section 3.3. In the current prototype, each snapshot is stored in a directory on HDFS. The name of the directory includes the progress value associated with the snapshot. Each reduce task runs independently, and at a different rate; once a reduce task has made sufficient progress, it writes a snapshot to a temporary directory on HDFS, and then atomically renames it to the appropriate snapshot directory.

Applications can consume snapshots by polling HDFS in a predictable location. An application knows that a given snapshot has been completed when every reduce task has written a file to the snapshot directory. Atomic rename is used to avoid applications mistakenly reading incomplete snapshot files (the same technique is used by stock Hadoop when writing the final output of each reduce task).

Note that if there are not enough free slots to allow all the reduce tasks in a job to be scheduled, snapshots will not be available for reduce tasks that are still waiting to be executed. The user can detect this situation (e.g. by checking for the expected number of files in the HDFS snapshot directory), so there is no risk of incorrect data, but the usefulness of online aggregation will be compromised. In the current prototype, we manually configured the cluster to avoid this scenario. The system could also be enhanced to avoid this pitfall entirely, e.g. by optionally waiting to execute an online aggregation job until there are enough reduce slots available.

### 4.1.1 Progress Metric

Hadoop provides support for monitoring the progress of task executions. As each map task executes, it is assigned a *progress score* in the range  $[0,1]$ , based on how much of its input the map task has consumed. We reused this feature to determine how much progress is represented by the current input to a reduce task, and hence to decide

when a new snapshot should be taken.

First, we modified the spill file format depicted in Figure 2 to include the map’s current progress score. When a partition in a spill file is sent to a reducer, the spill file’s progress score is also included. To compute the progress score for a snapshot, we take the average of the progress scores associated with each spill file used to produce the snapshot.

Note that it is possible that a map task might not have pipelined *any* output to a reduce task, either because the map task has not been scheduled yet (there are no free TaskTracker slots), or because the reduce task has been configured to only pipeline data from at most  $k$  map tasks concurrently. To account for this, we need to scale the progress metric to reflect the portion of the map tasks that a reduce task has pipelined data from: if a reducer is connected to  $\frac{1}{n}$  of the total number of map tasks in the job, we divide the average progress score by  $n$ .

This progress metric could easily be made more sophisticated: for example, an improved metric might include the selectivity ( $|output|/|input|$ ) of each map task, the statistical distribution of the map task’s output, and the effectiveness of each map task’s combine function, if any. Although we have found our simple progress metric to be sufficient for the experiments we describe below, this clearly represents an opportunity for future work.

#### 4.1.2 Evaluation

We found that our implementation of online aggregation produced an accurate estimate very quickly, and contributed relatively little overhead to job completion times. Our metric of accuracy for this experiment is post-hoc — we note the time at which the Top- $K$  words in the snapshot are the Top- $K$  in the final result.

Figure 6 contains two graphs that report the result of a job that searches for the top 5 words in a 550MB data set taken from Project Gutenberg. The correct Top-5 answers were available within about 21% of the total runtime. The 128MB block size used to load the data made a difference in the long-term execution behavior of this job. The file was broken up into 5 blocks, one of which (the tail) was a small block. The task assigned to this small block finishes almost immediately. The blocking job in the left graph shows this effect, since reduce tasks receive the output of this fast map early. Three of the remaining map tasks finish around the same time, with one map falling slightly behind. The right graph depicts pipelining with online aggregation. Although online aggregation does incur a performance penalty, it is relatively modest. Moreover, an accurate approximation of the answer is computed well before the final completion time of either job.

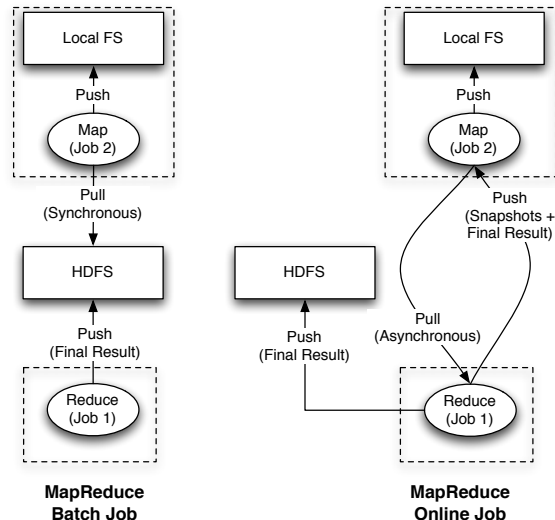


Figure 7: Hadoop dataflow for batch (left) and snapshot (right) processing of MapReduce jobs.

## 4.2 Multi-Job Online Aggregation

Online aggregation is particularly useful when applied to a long-running analysis task composed of multiple MapReduce jobs. As described in Section 3.2, our version of Hadoop allows the output of a reduce task to be sent directly to map tasks. This feature can be used to support online aggregation for a sequence of jobs.

Suppose that  $j_1$  and  $j_2$  are two MapReduce jobs, and  $j_2$  consumes the output of  $j_1$ . When  $j_1$ ’s reducers compute a snapshot to perform online aggregation, that snapshot is written to HDFS, and also sent directly to the map tasks of  $j_2$ . The map and reduce steps for  $j_2$  are then computed as normal, to produce a snapshot of  $j_2$ ’s output. This process can then be continued to support online aggregation for an arbitrarily long sequence of jobs.

Unfortunately, inter-job online aggregation has some drawbacks. First, the output of a reduce function is not “monotonic”: the output of a reduce function on the first 50% of the input data may not be obviously related to the output of the reduce function on the first 25%. Thus, as new snapshots are produced by  $j_1$ ,  $j_2$  must be recomputed from scratch using the new snapshot. As with inter-job pipelining (Section 3.2), this could be optimized for reduce functions that are declared to be distributive or algebraic aggregates [8].

To make inter-job online aggregation fault tolerant, we must consider three cases. Tasks that fail in  $j_1$  recover as described in Section 3.3. If a task in  $j_2$  fails, the system simply restarts the failed task. Since the snapshots produced by  $j_1$  are monotonic, the next snapshot received by the restarted task in  $j_2$  will have a higher progress score. To handle failures in  $j_1$ , tasks in  $j_2$  cache the most




Figure 6: Single job that finds the top 5 words in the 550MB Gutenberg data set. In the graph on the right, the vertical dashed line indicates when online aggregation produces an accurate result.

recent snapshot received by  $j_1$ , and replace it when they receive a new snapshot with a higher progress. If tasks from both jobs fail, a new task in  $j_2$  recovers the most recent snapshot from  $j_1$  that was stored HDFS and then wait for snapshots with a higher progress.

Figure 7 depicts an inter-job dataflow for batch mode (left) and a dataflow that supports online aggregation (right). The batch mode forces jobs to read the final output of other jobs from HDFS.

In addition to reading snapshots from HDFS, we support pipelining snapshots directly to the jobs that request them. This is supported via an asynchronous request call interface exported by each reduce task, through which map tasks from other jobs can request snapshots and even the final output. The final output is concurrently written to HDFS for fault tolerance, unless otherwise specified in the configuration of the job.

#### 4.2.1 Evaluation

To evaluate inter-job dataflow with online aggregation we wrote a Top- $K$  query using two MapReduce jobs and executed it on 5.5GB of Wikipedia article text. The first job performs a wordcount on the words contained in each article. A reducer from the first job will output a list of the Top- $K$  words observed in its partition. The key in this output is the word, and the value is the word count. Each map task in the subsequent job is assigned to an output from a single reduce task. The map function reverses the key-value order, and sends that result (sorted by count in descending order) to a single reduce task. The single reduce task merges the sorted lists from each mapper and returns the first  $K$  words.

Figure 8 reports the result of a Top-100 query over 5.5GB of Wikipedia article text. The left graph represents the blocking case as indicated by the idle period in the progress of the reduce tasks. The first map tasks finish around 100 seconds into the job, which is where the reduce tasks begin to make progress. The right graph

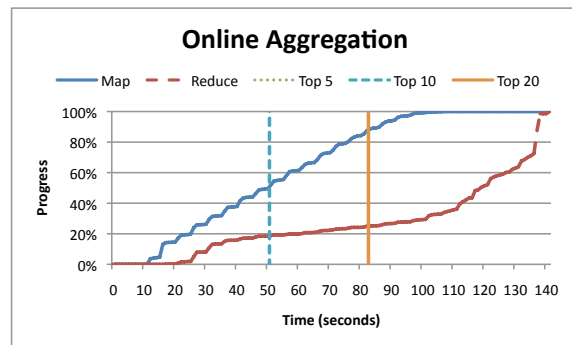


Figure 9: Multi-job Top-100 query over 5.5GB of Wikipedia article text. The vertical lines describe the increasing accuracy of the intermediate results produced by online aggregation.

shows a more balanced load. The reduce tasks being receiving mapper output almost immediately following the job execution, contributing to the early completion time of the job.

Figure 9 contains a graph of a Top-100 query over the same Wikipedia data set but this time executing the job with online aggregation. Although the final result for this job did not appear until nearly the end, we did observe the Top-5, 10, and 20 values at the times indicated in the graph. The job performed its online aggregation work as frequently as possible, which resulted in a performance hit in the final job completion time.

## 5 Continuous Queries

MapReduce is often used to analyze streams of constantly-arriving data, such as URL access logs [6] and system console logs [28]. Because of traditional constraints on MapReduce, this is done in large batches that can only provide periodic views of activity. This introduces significant latency into a data analysis process that ideally should run in near-real time. It is also potentially




Figure 8: Multi-job Top-100 query over 5.5GB of Wikipedia article text.

inefficient: each new MapReduce job does not have access to the computational state of the last analysis run, so this state must be recomputed from scratch. The programmer can manually save the state of each job and then reload it for the next analysis operation, but this is labor-intensive.

Our pipelined version of Hadoop allows an alternative architecture: MapReduce jobs that run *continuously*, accepting new data as it becomes available and analyzing it immediately. This allows for near-real-time analysis of data streams, and thus allows the MapReduce programming model to be applied to domains such as environment monitoring and real-time fraud detection.

In this section, we describe how HOP supports continuous MapReduce jobs, and how we used this feature to implement a rudimentary cluster monitoring tool.

## 5.1 Continuous MapReduce Jobs

A bare-bones implementation of continuous MapReduce jobs is easy to implement using pipelining. No changes are needed to implement continuous map tasks: map output is already delivered to the appropriate reduce task shortly after it is generated. We added an optional “flush” API that allows map functions to force their current output to reduce tasks. When a reduce task is unable to accept such data, the mapper framework stores it locally and sends it at a later time. With proper scheduling of reducers, this API allows a map task to ensure that an output record is promptly sent to the appropriate reducer.

To support continuous reduce tasks, the user-defined reduce function must be periodically invoked on the map output available at that reducer. Applications will have different requirements for how frequently the reduce function should be invoked: possible choices include periods based on wall-clock time, logical time (e.g. the value of a field in the map task output), and the number of input rows delivered to the reducer. The output of the reduce function can be written to HDFS, as in our implementation of online aggregation. However, other

choices are possible: our prototype system monitoring application (described below) sends an alert via email if an anomalous situation is detected.

In our current implementation, the number of map and reduce tasks is fixed, and must be configured by the user. This is clearly problematic: manual configuration is error-prone, and many stream processing applications exhibit “bursty” traffic patterns, in which peak load far exceeds average load. In the future, we plan to add support for elastic scaleup/scaledown of map and reduce tasks, in response to variations in load.

### 5.1.1 Fault Tolerance

In the checkpoint/restart fault-tolerance model used by Hadoop, mappers retain their output until the end of the job to facilitate fast recovery from reducer failures. In a continuous query context, this is infeasible, since mapper history is in principle infinite. However, many continuous reduce functions (e.g., 30-second moving average) only depend on a suffix of the history of the map stream. This common case can be supported easily, by extending the JobTracker interface to capture a rolling notion of reducer consumption. Map-side spill files are maintained in a ring buffer with unique IDs for spill files over time. When a reducer commits an output to HDFS, it informs the JobTracker about the *run* of map output records it no longer needs, identifying the run by spill file IDs and offsets within those files. The JobTracker then can tell mappers to delete the appropriate data.

In principle, complex reducers may depend on very long (or infinite) histories of map records to accurately reconstruct their internal state. In that case, deleting spill files from the map-side ring buffer will result in potentially inaccurate recovery after faults. Such scenarios can be handled by having reducers checkpoint internal state to HDFS, along with markers for the mapper offsets at which the internal state was checkpointed. The MapReduce framework can be extended with APIs to help with state serialization and offset management, but




Figure 10: Pages swapped in and out over time on the thrashing host. The vertical line indicates the time at which the alert was sent by the monitoring system.

it still presents a programming burden on the user to correctly identify the sensitive internal state. That burden can be avoided by more heavyweight process-pair techniques for fault tolerance, but those are quite complex and use significant resources [23]. In our work to date we have focused on cases where reducers can be recovered from a reasonable-sized history at the mappers, favoring minor extensions to the simple fault-tolerance approach used in Hadoop.

## 5.2 Prototype Monitoring System

Our monitoring system is composed of *agents* that run on each monitored machine and record statistics of interest (e.g. load average, I/O operations per second, etc.). Each agent is implemented as a continuous map task: rather than reading from HDFS, the map task instead reads from various continuous system-local data streams (e.g. `/proc`).

Each agent forwards statistics to an *aggregator* that is implemented as a continuous reduce task. The aggregator records how agent-local statistics evolve over time (e.g. by computing windowed-averages), and compares statistics between agents to detect anomalous behavior. Each aggregator only monitors the agents that report to it, but might also report statistical summaries to another “upstream” aggregator. For example, the system might be configured to have an aggregator for each rack, and then a second level of aggregators that compare statistics between racks to analyze datacenter-wide behavior.

## 5.3 Evaluation

To validate our prototype system monitoring tool, we constructed a scenario in which one member of a MapReduce cluster begins thrashing during the execution of a job. Our goal was to test how quickly our monitoring system would detect this behavior. The basic mechanism is similar to an alert system one of the

authors implemented at an Internet search company.

We used a simple load metric (a linear combination of CPU utilization, paging, and swap activity). The continuous reduce function maintains windows over samples of this metric: at regular intervals, it compares the 20 second moving average of the load metric for each host to the 120 second moving average of all the hosts in the cluster *except* that host. If the given host’s load metric is more than two standard deviations above the global average, it is considered an outlier and a tentative alert is issued. To dampen false positives in “bursty” load scenarios, we do not issue an alert until we have received 10 tentative alerts within a time window.

We deployed this system on an EC2 cluster consisting of 7 “large” nodes (large nodes were chosen because EC2 allocates an entire physical host machine to them). We ran a wordcount job on the 5.5GB Wikipedia data set, using 5 map tasks and 2 reduce tasks (1 task per host). After the job had been running for about 10 seconds, we selected a node running a task and launched a program that induced thrashing.

We report detection latency in Figure 10. On the thrashing host, we ran `vmstat` to monitor the number of pages swapped over time. The vertical bar indicates the time at which the monitoring tool fired a (non-tentative) alert. The thrashing host was detected very rapidly—notably faster than the 5-second TaskTracker-JobTracker heartbeat cycle that is used to detect straggler tasks in stock Hadoop. We envision using these alerts to do early detection of stragglers within a MapReduce job. This is a case of HOP making scheduling decisions for a MapReduce job by running a secondary continuous MapReduce monitoring query. Compared to out-of-band monitoring tools, this economy of mechanism—reusing the MapReduce infrastructure for reflective monitoring—has benefits in software maintenance and system management.

## 6 Related Work

The work in this paper relates to literature on parallel dataflow frameworks, online aggregation, and continuous query processing.

### 6.1 Parallel Dataflow

Dean and Ghemawat’s paper on Google’s MapReduce [6] has become a standard reference, and forms the basis of the open-source Hadoop implementation. As noted in Section 1, the Google MapReduce design targets very large clusters where the probability of worker failure or slowdown is high. This led to their elegant checkpoint/restart approach to fault tolerance, and (as a corollary) their lack of pipelining. Our work extends the

Google design to accommodate pipelining without significant modification to their core programming model or fault tolerance mechanisms.

*Dryad* [12] is a data-parallel programming model and runtime that is often compared to MapReduce, supporting a more general model of acyclic dataflow graphs. Like MapReduce, Dryad puts disk materialization steps between dataflow stages by default, breaking pipelines. The Dryad paper describes support for optionally “encapsulating” multiple stages of special asynchronous stages into a single process so they can pipeline, but this requires a more complicated programming interface. The Dryad paper explicitly mentions that the system is targeted at batch processing, and not at scenarios like continuous queries.

It has been noted that parallel database systems have long provided partitioned dataflow frameworks [20], and recent commercial databases have begun to offer MapReduce programming models on top of those frameworks [5, 9]. Most of the parallel database systems can provide pipelined execution akin to our work here, but they use a more tightly coupled iterator and *Exchange* model that keeps producers and consumers rate-matched via queues, spreading the work of each dataflow stage across all nodes in the cluster [7]. This provides less scheduling flexibility than Hadoop’s MapReduce model, and typically offers no tolerance to mid-query worker faults. Aster Data advertises support for mid-query fault-tolerance, though their technical report on the MapReduce implementation does not explain how fault tolerance is achieved [5].

Logothetis and Yocum describe a MapReduce interface over a continuous query system called *Mortar* that is similar in some ways to our work [15]. Like HOP, their mappers push data to reducers in a pipelined fashion. They focus on specific issues in efficient stream query processing, including minimization of work for aggregates in overlapping windows via special reducer APIs. They are not built on Hadoop, and explicitly sidestep issues in fault tolerance.

*Hadoop Streaming* is an interface that allows Map and Reduce functions to be expressed as UNIX shell command lines. It does not stream data through Map and Reduce phases in a pipelined fashion.

## 6.2 Online Aggregation

Online aggregation was originally proposed in the context of simple single-table SQL queries involving “Group By” aggregations, a workload quite similar to MapReduce [11]. The focus of the initial work was on providing not only “early returns” to these SQL queries, but also statistically robust estimators and confidence interval metrics for the final result based on random sam-

pling. These statistical matters do not generalize to arbitrary MapReduce jobs, though our framework can support those that have been developed. Subsequently, online aggregation was extended to handle join queries (via the *Ripple Join* method), and the *CONTROL* project generalized the idea of online query processing to provide interactivity for data cleaning, data mining, and data visualization tasks [10]. That work was targeted at single-processor systems. Luo et al. developed a partitioned-parallel variant of Ripple Join, without statistical guarantees on approximate answers [16].

In recent years, this topic has seen renewed interest, starting with Jermaine et al.’s work on the *DBO* system [14]. That effort includes more disk-conscious online join algorithms, as well as techniques for maintaining randomly-shuffled files to remove any potential for statistical bias in scans [13]. Wu et al. describe a system for peer-to-peer online aggregation in a distributed hash table context [27]. The open programmability and fault-tolerance of MapReduce are not addressed significantly in prior work on online aggregation.

## 6.3 Continuous Queries

In the last decade there was a great deal of work in the database research community on the topic of continuous queries over data streams, including systems such as Borealis [1], STREAM [17], and Telegraph [4]. Of these, Borealis and Telegraph [23] studied fault tolerance and load balancing across machines. In the Borealis context this was done for pipelined dataflows, but without partitioned parallelism: each stage (“operator”) of the pipeline runs serially on a different machine in the wide area, and fault tolerance deals with failures of entire operators [3]. SBON [21] is an overlay network that can be integrated with Borealis, which handles “operator placement” optimizations for these wide-area pipelined dataflows.

Telegraph’s *FLuX* operator [24, 23] is the only work to our knowledge that addresses mid-stream fault-tolerance for dataflows that are both pipelined and partitioned in the style of HOP. FLuX (“Fault-tolerant, Load-balanced eXchange”) is a dataflow operator that encapsulates the shuffling done between stages such as map and reduce. It provides load-balancing interfaces that can migrate operator state (e.g., reducer state) between nodes, while handling scheduling policy and changes to data-routing policies [24]. For fault tolerance, FLuX develops a solution based on process pairs [23], which work redundantly to ensure that operator state is always being maintained live on multiple nodes. This removes any burden on the continuous query programmer of the sort we describe in Section 5. On the other hand, the FLuX protocol is far more complex and resource-intensive than our pipelined adap-

tation of Google’s checkpoint/restart tolerance model.

## 7 Conclusion and Future Work

MapReduce has proven to be a popular model for large-scale parallel programming. Our Hadoop Online Prototype extends the applicability of the model to pipelining behaviors, while preserving the simple programming model and fault tolerance of a full-featured MapReduce framework. This provides significant new functionality, including “early returns” on long-running jobs via online aggregation, and continuous queries over streaming data. We also demonstrate benefits for batch processing: by pipelining both within and across jobs, HOP can reduce the time to job completion.

In considering future work, scheduling is a topic that arises immediately. Stock Hadoop already has many degrees of freedom in scheduling batch tasks across machines and time, and the introduction of pipelining in HOP only increases this design space. First, pipeline parallelism is a new option for improving performance of MapReduce jobs, but needs to be integrated intelligently with both intra-task partition parallelism, and speculative redundant execution for “straggler” handling. Second, the ability to schedule deep pipelines with direct communication between reduces and maps (bypassing the distributed file system) opens up new opportunities and challenges in carefully co-locating tasks from different jobs, to avoid communication when possible.

Olston and colleagues have noted that MapReduce systems—unlike traditional databases—require “model-light” optimization approaches that gather and react to performance information during runtime [18]. The continuous query facilities of HOP enable powerful introspective programming interfaces for this: a full-featured MapReduce interface can be used to script performance monitoring tasks that gather system-wide information in near-real-time, enabling tight feedback loops for scheduling and dataflow optimization. This is a topic we plan to explore, including opportunistic methods to do monitoring work with minimal interference to outstanding jobs, as well as dynamic approaches to continuous optimization in the spirit of earlier work like Eddies [2] and Flux [24].

As a more long-term agenda, we want to explore using MapReduce-style programming for even more interactive applications. As a first step, we hope to revisit interactive data processing in the spirit of the CONTROL work [10], with an eye toward improved scalability via parallelism. More aggressively, we are considering the idea of bridging the gap between MapReduce dataflow programming and lightweight event-flow programming models like SEDA [26]. Our HOP implementation’s roots in Hadoop make it unlikely to compete

with something like SEDA in terms of raw performance. But it would be interesting to translate ideas across these two traditionally separate programming models, perhaps with an eye toward building a new and more general-purpose framework for programming in architectures like cloud computing and many-core.

## Acknowledgments

We would like to thank Kuang Chen, Akshay Krishnamurthy, and Mosharaf Chowdhury for their helpful comments. This research is supported by grants from the National Science Foundation and the Natural Sciences and Engineering Research Council of Canada.

## References

- [1] ABADI, D. J., AHMAD, Y., BALAZINSKA, M., CETINTEMEL, U., CHERNIACK, M., HWANG, J.-H., LINDNER, W., MASKEY, A. S., RASIN, A., RYVKINA, E., TATBUL, N., XING, Y., AND ZDONIK, S. The design of the Borealis stream processing engine. In *CIDR* (2005).
- [2] AVNUR, R., AND HELLERSTEIN, J. M. Eddies: Continuously adaptive query processing. In *SIGMOD* (2000), pp. 261–272.
- [3] BALAZINSKA, M., BALAKRISHNAN, H., MADDEN, S., AND STONEBRAKER, M. Fault-tolerance in the Borealis distributed stream processing system. In *SIGMOD* (2005).
- [4] CHANDRASEKARAN, S., COOPER, O., DESHPANDE, A., FRANKLIN, M. J., HELLERSTEIN, J. M., HONG, W., KRISHNAMURTHY, S., MADDEN, S. R., RAMAN, V., REISS, F., AND SHAH, M. A. TelegraphCQ: Continuous dataflow processing for an uncertain world. In *CIDR* (2003).
- [5] CIESLEWICZ, J., FRIEDMAN, E., AND PAWLOWSKI, P. SQL/MapReduce: A practical approach to self-describing, polymorphic, and parallelizable user-defined functions. In *VLDB* (2009).
- [6] DEAN, J., AND GHEMAWAT, S. MapReduce: Simplified data processing on large clusters. In *OSDI* (2004), pp. 137–150.
- [7] GRAEFE, G. Encapsulation of parallelism in the Volcano query processing system. In *SIGMOD* (1990).
- [8] GRAY, J., CHAUDHURI, S., BOSWORTH, A., LAYMAN, A., REICHART, D., VENKATRAO, M., PELLOW, F., AND PIRAHESH, H. Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub totals. *Data Min. Knowl. Discov. J.*, 1 (1997), 29–53.
- [9] Greenplum: A unified engine for RDBMS and MapReduce, Oct. 2008. Downloaded from <http://www.greenplum.com/download.php?alias=register-map-reduce&file=Greenplum-MapReduce-Whitepaper.pdf>.
- [10] HELLERSTEIN, J. M., AVNUR, R., CHOU, A., HIDBER, C., OLSTON, C., RAMAN, V., ROTH, T., AND HAAS, P. J. Interactive data analysis with CONTROL. *IEEE Computer* 32, 8 (Aug. 1999).
- [11] HELLERSTEIN, J. M., HAAS, P. J., AND WANG, H. J. Online aggregation. In *SIGMOD* (1997), pp. 171–182.
- [12] ISARD, M., BUDI, M., YU, Y., BIRRELL, A., AND FETTERLY, D. Dryad: Distributed data-parallel programs from sequential building blocks. In *EuroSys* (2007).

- [13] JERMAINE, C. Online random shuffling of large database tables. *IEEE Trans. Knowl. Data Eng.* 19, 1 (2007), 73–84.
- [14] JERMAINE, C., ARUMUGAM, S., POL, A., AND DOBRA, A. Scalable approximate query processing with the DBO engine. In *SIGMOD* (2007).
- [15] LOGOTHETIS, D., AND YOCUM, K. Ad-hoc data processing in the cloud (demonstration). *Proc. VLDB Endow.* 1, 2 (2008).
- [16] LUO, G., ELLMANN, C. J., HAAS, P. J., AND NAUGHTON, J. F. A scalable hash ripple join algorithm. In *SIGMOD* (2002).
- [17] MOTWANI, R., WIDOM, J., ARASU, A., BABCOCK, B., BABU, S., DATAR, M., MANKU, G., OLSTON, C., ROSENSTEIN, J., AND VARMA, R. Query processing, resource management, and approximation in a data stream management system. In *CIDR* (2003).
- [18] OLSTON, C., REED, B., SILBERSTEIN, A., AND SRIVASTAVA, U. Automatic optimization of parallel dataflow programs. In *USENIX Technical Conference* (2008).
- [19] OLSTON, C., REED, B., SRIVASTAVA, U., KUMAR, R., AND TOMKINS, A. Pig Latin: a not-so-foreign language for data processing. In *SIGMOD* (2008).
- [20] PAVLO, A., PAULSON, E., RASIN, A., ABADI, D. J., DEWITT, D. J., MADDEN, S., AND STONEBRAKER, M. A comparison of approaches to large-scale data analysis. In *SIGMOD* (2009).
- [21] PIETZUCH, P., LEDLIE, J., SHNEIDMAN, J., ROUSSOPOULOS, M., WELSH, M., AND SELTZER, M. Network-aware operator placement for stream-processing systems. In *ICDE* (2006).
- [22] PIKE, R., DORWARD, S., GRIESEMER, R., AND QUINLAN, S. Interpreting the data: Parallel analysis with Sawzall. *Scientific Programming* 13, 4 (2005), 277–298.
- [23] SHAH, M. A., HELLERSTEIN, J. M., AND BREWER, E. A. Highly-available, fault-tolerant, parallel dataflows. In *SIGMOD* (2004).
- [24] SHAH, M. A., HELLERSTEIN, J. M., CHANDRASEKARAN, S., AND FRANKLIN, M. J. Flux: An adaptive partitioning operator for continuous query systems. In *ICDE* (2003).
- [25] THUSOO, A., SARMA, J. S., JAIN, N., SHAO, Z., CHAKKA, P., ANTHONY, S., LIU, H., WYCKOFF, P., AND MURTHY, R. Hive — a warehousing solution over a map-reduce framework. In *VLDB* (2009).
- [26] WELSH, M., CULLER, D., AND BREWER, E. SEDA: An architecture for well-conditioned, scalable internet services. In *SOSP* (2001).
- [27] WU, S., JIANG, S., OOI, B. C., AND TAN, K.-L. Distributed online aggregation. In *VLDB* (2009).
- [28] XU, W., HUANG, L., FOX, A., PATTERSON, D., AND JORDAN, M. I. Detecting large-scale system problems by mining console logs. In *SOSP* (2009).