

# HAMAKE: A DATAFLOW APPROACH TO DATA PROCESSING IN HADOOP

Vadim Zaliva<sup>1</sup> and Vladimir Orlov<sup>2</sup>

<sup>1</sup>*Codeminders, Saratoga, CA, U.S.A.*

<sup>2</sup>*Codeminders, Kiev, Ukraine*

Keywords: Hadoop, Hamake, Dataflow.

Abstract: Most non-trivial data processing scenarios using Hadoop typically involve launching more than one MapReduce job. Usually, such processing is data-driven with the data funneled through a sequence of jobs. The processing model could be expressed in terms of dataflow programming, represented as a directed graph with datasets as vertices. Using *fuzzy timestamps* as a way to detect which dataset needs to be updated, we can calculate a sequence in which Hadoop jobs should be launched to bring all datasets up to date. Incremental data processing and parallel job execution fit well into this approach. These ideas inspired the creation of the hamake utility. We attempted to emphasize data allowing the developer to formulate the problem as a data flow, in contrast to the workflow approach commonly used. Hamake language uses just two data flow operators: *fold* and *foreach*, providing a clear processing model similar to MapReduce, but on a dataset level.

## 1 MOTIVATION AND BACKGROUND

Hadoop (Bialecki et al., 2012) is a popular open-source implementation of MapReduce, a data processing model introduced by Google (Dean and Ghemawat, 2008).

Hadoop is typically used to process large amounts of data through a series of relatively simple operations. Usually Hadoop jobs are I/O-bound (Weil, 2012; Gangadhar, 2012), and execution of even trivial operations on a large dataset could take significant system resources. This makes incremental processing especially important. Our initial inspiration was the Unix *make* utility. While applying some of the ideas implemented by *make* to Hadoop, we took the opportunity to generalize the processing model in terms of dataflow programming.

**Hamake** was developed to address the problem of incremental processing of large data sets in a collaborative filtering project.

We've striven to create an easy to use utility that developers can start using right away without complex installation or extensive learning curve.

**Hamake** is open source and is distributed under Apache License v2.0. The project is hosted at Google Code at the following URL: <http://code.google.com/p/hamake/>.

## 2 PROCESSING MODEL

**Hamake** operates on *files* residing on a local or distributed file system accessible from the Hadoop job. Each file has a timestamp reflecting the date and time of its last modification. A file system directory or folder is also a file with its own timestamp. A *Data Transformation Rule (DTR)* defines an operation which takes files as input and produces other files as output.

If file *A* is listed as input of a DTR, and file *B* is listed as output of the same DTR, it is said that "*B depends on A.*" **Hamake** uses file time stamps for dependency up-to-date checks. DTR output is said to be *up to date* if the minimum time stamp on all outputs is greater than or equal to the maximum timestamp on all inputs. For the sake of convenience, a user could arrange groups of files and folders into a *fileset* which could later be referenced as the DTR's input or output.

**Hamake** uses *fuzzy timestamps*<sup>1</sup> which can be compared, allowing for a slight margin of error. The "fuzziness" is controlled by a tolerance of  $\sigma$ . Timestamp *a* is considered to be older than timestamp *b* if  $(b - a) > \sigma$ . Setting  $\sigma = 0$  gives us a non-fuzzy, strict timestamp comparison.

**Hamake** attempts to ensure that all outputs from

<sup>1</sup>The current stable version of **hamake** uses exact (non-fuzzy) timestamps.

a DTR are up to date<sup>2</sup> To do so, it builds a *dependency graph* with DTRs as edges and individual files or filesets as vertices. Below, we show that this graph is guaranteed to be a *Directed Acyclic Graph* (DAG).

After building a *dependency graph*, a graph reduction algorithm (shown in Figure 1) is executed. Step 1 uses Kahn's algorithm (Kahn, 1962) of topological ordering. In step 6, when the completed DTR is removed from the dependency graph, all edges pointing to it from other DTRs are also removed.

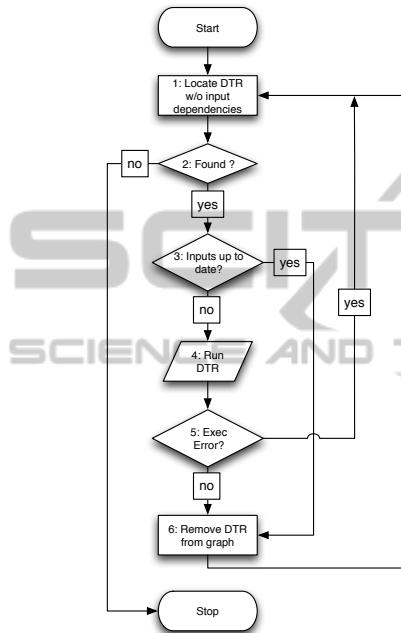


Figure 1: **Hamake** dependency graph reduction algorithm.

The algorithm allows for parallelism. If more than one DTR without input dependencies is found during step 1, the subsequent steps 2-6 can be executed in parallel for each discovered DTR.

It should be noted that if DTR execution has failed, **hamake** can and will continue to process other DTRs which do not depend directly or indirectly on the results of this DTR. This permits the user to fix problems later and re-run **hamake**, without the need to re-process all data.

Cyclic dependencies must be avoided, because a dataflow containing such dependencies is not guaranteed to terminate. Implicit checks are performed during the reading of DAG definitions and the building of the dependency graph. If a cycle is detected, it is reported as an error. Thus the dependency graph used by **hamake** is assured to be a *directed acyclic graph*.

<sup>2</sup>Because **hamake** has no way to update them, it does not attempt to ensure that files are up to date, unless they are listed as one of a DTR's outputs.

However, **hamake** supports a limited scenario of iterative processing with a feature called *generations*. Each input or output file can be marked with a *generation* attribute. Any two files referencing the same path in the file system while having different generations are represented as two distinct vertices in the dependency graph. This permits resolution of cyclic dependencies within the context of a single **hamake** execution.

One useful consequence of **hamake** dataflow being a DAG is that for each vertex we can calculate the list of vertices it depends on directly and indirectly using simple *transitive closure*. This allows us to easily estimate the part of a dataflow graph being affected by updating one or more files, which could be especially useful for datasets where the cost of re-calculation is potentially high due to data size or computational complexity.

**Hamake** is driven by dataflow description, expressed in a simple XML-based language. The full syntax is described in (Orlov and Bondar, 2012). The two main elements, *fold* and *foreach*, correspond to two types of DTRs. Each element has input, output, and processing instructions. The execution of processing instructions brings the DTR output up to date.

*Fold* implies a many-to-one dependency between input and output. In other words, the output depends on the entirety of the input, and if any of the inputs have been changed, the outputs need to be updated. *Foreach* implies a one-to-one dependency where for each file in an input set there is a corresponding file in an output set, each updated independently.

**Hamake** dataflow language has declarative semantics making it easy to implement various dataflow analysis and optimization algorithms in the future. Examples of such algorithms include: merging dataflows, further execution parallelization, and analysis and estimation of dataflow complexity.

### 3 PARALLEL EXECUTION

While determining the sequence and launching of Hadoop jobs required to bring all datasets up-to-date, **hamake** attempts to perform all required computations in the shortest possible time. To achieve this, **hamake** aims for maximal cluster utilization, running as many Hadoop jobs in parallel as cluster capacity permits.

There are three main factors that drive job scheduling logic: file timestamps, dependencies, and cluster computational capacity. On the highest level, DTR dependencies determine the sequence of jobs to be launched.

In the case of *fold* DTR, a single Hadoop job, PIG script or shell command, may be launched, and hence there is no opportunity for parallel execution. In the example shown in Figure 2, since fileset *B* depends on all files in fileset *A*, a single job associated with *fold* DTR will be executed.




Figure 2: Decomposition of *fold* DTR.

A *foreach* DTR works by mapping individual files in fileset *A* to files in fileset *B*. Assuming that fileset *A* consists of 3 files:  $a_1, a_2, a_3$ , the dependency graph could be represented as shown in Figure 3. In this case, we have an opportunity to execute the three jobs in parallel.




Figure 3: Decomposition of *foreach* DTR.

The Hadoop cluster capacity is defined in terms of the number of *map slots* and *reduce slots*. When a DTR launches a Hadoop job, either directly as defined by *mapreduce* processing instruction or via PIG script, a single job will spawn one or more *mapper* or *reducer* tasks, each taking one respective slot. The number of mappers and reducers launched depends on many factors, such as the size of the HDFS block, Hadoop cluster settings, and individual job settings. In general, **hamake** has neither visibility of nor control over most of these factors, so it does not currently deal with individual tasks. Thus **hamake** parallel execution logic is controlled by a command line option specifying how many jobs it may run in parallel.

## 4 EXAMPLE

In a large, online library, the **hamake** utility can be used to automate searches for duplicates within a corpus of millions of digital text documents. Documents with slight differences due to OCR errors, typos, differences in formatting, or added material such as a foreword or publishers note can be found and reported.

To illustrate **hamake** usage, consider the simple approach of using the *Canopy clustering algorithm* (McCallum et al., 2000) and a *vector space model* (Manning et al., 2008) based on word frequencies. The implementation could be split into a series of steps, each implemented as *MapReduce job*:

**ExtractText** Extract a plain text from native document format (e.g. PDF).

**Tokenize** Split plain text into tokens which roughly correspond to words. Deal with hyphens, compound words, accents, and diacritics, as well as case-folding, stemming, or lemmatization, resulting in a list of normalized tokens.

**FilterStopwords** Filter out *stopwords*, like *a*, *the*, and *are*.

**CalculateTF** Calculate a feature vector of term frequencies for each document.

**FindSimilar** Run *Canopy clustering algorithm* to group similar documents into clusters using *co-sine distance* as a fast approximate distance metric.

**OutputResult** Output document names, which are found in clusters with more than one element.

Each of the six MapReduce jobs produces an output file which depends on its input. For each document, these jobs must be invoked sequentially, as the output of one task is used as input of the next. Additionally, there is a configuration file containing a list of stop words, and some task outputs depend on this file content. These dependencies could be represented by a DAG, as shown in Figure 4, with vertices representing documents and jobs assigned to edges. The XML file describing this dataflow in **hamake** language is shown as Listing 1.

Listing 1: *hamakefile*, describing process for detecting duplicate documents.

```

1<?xml version="1.0" encoding="UTF-8"?>
2<project name="FindSimilarBooks">
3
4  <property name="lib" value="/lib/">
5
6  <fileset id="input" path="/doc" mask="*.pdf"/>
7  <file id="output" path="/result.txt"/>
8
9  <foreach name="ExtractText">
10   <input>

```

```

11 <include idref="input"/>
12 </input>
13 <output>
14 <file id="plainText" path="/txt/${foreach:filename}"/>
15 </output>
16 <mapreduce jar="${lib}/hadoopJobs.job"
17 main="TextExtractor">
18 <parameter>
19 <literal value="${foreach:path}"/>
20 </parameter>
21 <parameter>
22 <reference idref="plainText"/>
23 </parameter>
24 </mapreduce>
25 </foreach>
26
27 <foreach name="Tokenize">
28 <input>
29 <file id="plainText" path="/txt"/>
30 </input>
31 <output>
32 <file id="tokens" path="/tokens/${foreach:filename}"/>
33 </output>
34 <mapreduce jar="${lib}/hadoopJobs.job"
35 main="Tokenizer">
36 ...
37 </mapreduce>
38 </foreach>
39
40 <foreach name="FilterStopWords">
41 <input>
42 <file id="stopWords" path="/stopwords.txt"/>
43 <file id="tokens" path="/tokens"/>
44 </input>
45 <output>
46 <file id="terms" path="/terms/${foreach:filename}"/>
47 </output>
48 <mapreduce jar="${lib}/hadoopJobs.job"
49 main="Tokenizer">
50 ...
51 </mapreduce>
52 </foreach>
53
54 <foreach name="CalculateTF">
55 <input>
56 <file id="terms" path="/terms"/>
57 </input>
58 <output>
59 <file id="TFVector" path="/TF"/>
60 </output>
61 <mapreduce jar="${dist}/hadoopJobs.job"
62 main="CalculateTF">
63 ...
64 </mapreduce>
65 </foreach>
66
67 <fold name="FindSimilar">
68 <input>
69 <file id="TFVector" path="/TF"/>
70 </input>
71 <output>
72 <include idref="clustersList" path="/clusters"/>
73 </output>
74 <mapreduce jar="${lib}/hadoopJobs.job"
75 main="Canopy">
76 ...
77 </mapreduce>
78 </fold>
79
80 <fold name="OutputResult">
81 <input>
82 <file id="clustersList" path="/clusters"/>
83 </input>
84 <output>
85 <include idref="output"/>
86 </output>
87 <mapreduce jar="${lib}/hadoopJobs.job"
88 main="OutputSimilarBooks">
89 ...
90 </mapreduce>
91 </fold>
92</project>

```

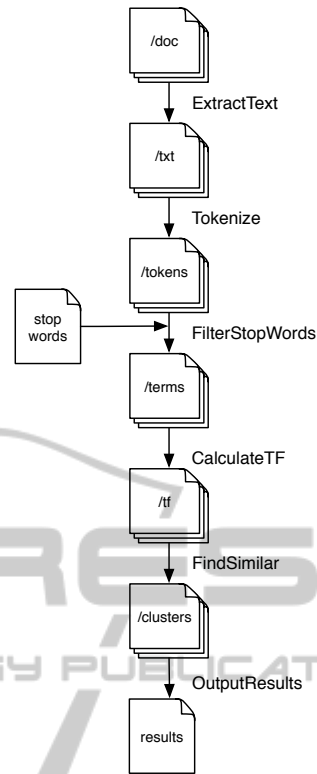


Figure 4: Directed acyclic graph of a data flow for duplicate document detection.

job which performs the actual text extraction is defined using the *mapreduce* element. It will be invoked by **hamake** for each unsatisfied dependency. The job takes two parameters, defined with *parameter* elements - a path to an original document as the input and a path to a file where the plain text version will be written. The remaining five DTRs are defined in a similar manner.

**Hamake**, when launched with this XML dataflow definition, will execute a graph reduction algorithm, as shown in Figure 1, and will find the first DTR to process. In our example, this is *ExtractPlainText*. First, **Hamake** will launch the corresponding Hadoop job and immediately following, execute DTRs which depend on the output of this DTR, and so on until all output files are up to date. As a result of this data flow, a file named *results.txt* with a list of similar documents will be generated.

This data flow could be used for incremental processing.

When new documents are added, **hamake** will refrain from running the following DTRs: *ExtractText*, *Tokenize*, *FilterStopWords*, and *CalculateTF* for previously processed documents. However, it will run those DTRs for newly added documents and then, re-run *FindSimilar* and *OutputResults*.

The first DTR (lines 10-25) converts a document from a native format such as PDF to plain text. The input of the DTR is a reference to the */doc* folder, and the output is the */txt* folder. The *foreach* DTR establishes one-to-one dependencies between files with identical names in these two folders. The Hadoop

If the list of stop words has been changed, **hamake** will re-run only *FilterStopWords*, *CalculateTF*, *FindSimilar*, and *OutputResults*.

## 5 RELATED WORK

Several workflow engines exist for Hadoop, such as Oozie (Yahoo!, 2012), Azkaban (Linkedin, 2012), Cascading (Wensel, 2012), and Nova (Olston et al., 2011). Although all of these products could be used to solve similar problems, they differ significantly in design, philosophy, target user profile, and usage scenarios limiting the usefulness of a simple, feature-wise comparison.

The most significant difference between these engines and **hamake** lies in the *workflow* vs. *dataflow* approach. All of them use the former, explicitly specifying dependencies between jobs. **Hamake**, in contrast, uses dependencies between datasets to derive workflow. Both approaches have their advantages, but for some problems, the dataflow representation as used by **hamake** is more natural.

Kangaroo (Zhang et al., 2011) is using similar data-flow DAG processing model, but not integrated with Hadoop.

## 6 FUTURE DIRECTIONS

One possible **hamake** improvement may be better integration with Hadoop schedulers. For example, if *Capacity Scheduler* or *Fair Scheduler* is used, it would be useful for **hamake** to take information about scheduler *pools* or *queues* capacity into account in its job scheduling algorithm.

More granular control over parallelism could be achieved if the **hamake** internal dependency graph for *foreach* DTR contained individual files rather than just filesets. For example, consider a dataflow consisting of three filesets *A*, *B*, *C*, and two *foreach* DTR's:  $D_1$ , mapping *A* to *B*, and  $D_2$ , mapping *B* to *C*. File-level dependencies would allow some jobs to run from  $D_2$  without waiting for all jobs in  $D_1$  to complete.

Another potential area of future extension is the **hamake** dependency mechanism. The current implementation uses a fairly simple timestamp comparison to check whether dependency is satisfied. This could be generalized, allowing the user to specify custom dependency check predicates, implemented either as plugins, scripts (in some embedded scripting languages), or external programs. This would allow for decisions based not only on file meta data, such as the timestamp, but also on its contents.

Several **hamake** users have requested support for iterative computations with a termination condition. Possible use-cases include fixed-point computations and clustering or iterative regression algorithms. Presently, to embed this kind of algorithm into the **hamake** dataflow, it requires the use of the *generations* feature combined with external automation, which invokes **hamake** repeatedly until a certain exit condition is satisfied. **Hamake** users could certainly benefit from native support for this kind of dataflow.

## REFERENCES

- Bialecki, A., Cafarella, M., Cutting, D., and O'Malley, O. (2005). Retrieved Feburay 06, 2012). Hadoop: a framework for running applications on large clusters built of commodity hardware.
- Dean, J. and Ghemawat, S. (2008). Map Reduce: Simplified data processing on large clusters. *Communications of the ACM-Association for Computing Machinery-CACM*, 51(1):107–114.
- Gangadhar, M. (2010). Retrieved Feburay 06, 2012). Benchmarking and optimizing hadoop. <http://www.slideshare.net/ydn/hadoop-summit-2010-benchmarking-and-optimizing-hadoop>.
- Kahn, A. (1962). Topological sorting of large networks. *Communications of the ACM*, 5(11):558–562.
- Linkedin (2010). Retrieved Feburay 08, 2012). Azkaban: Simple hadoop workflow. <http://sna-projects.com/azkaban/>.
- Manning, C., Raghavan, P., and Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press.
- McCallum, A., Nigam, K., and Ungar, L. H. (2000). Efficient Clustering of High-Dimensional Data Sets with Application to Reference Matching. *KDD '00*.
- Olston, C., Chiou, G., Chitnis, L., Liu, F., Han, Y., Larson, M., Neumann, A., Rao, V., Sankarasubramanian, V., Seth, S., et al. (2011). Nova: continuous pig/hadoop workflows. In *Proceedings of the 2011 international conference on Management of data*, pages 1081–1090. ACM.
- Orlov, V. and Bondar, A. (2011). Retrieved Feburay 06, 2012). Hamake syntax reference. <http://code.google.com/p/hamake/wiki/HamakeFileSyntaxReference>.
- Weil, K. (2010). Retrieved Feburay 06, 2012). Hadoop at twitter. <http://engineering.twitter.com/2010/04/hadoop-at-twitter.html>.
- Wensel, C. (2010). Retrieved Feburay 08, 2012). Cascading. <http://www.cascading.org/>.
- Yahoo! (2010). Retrieved Feburay 08, 2012). Oozie: Workflow engine for hadoop. <http://yahoo.github.com/oozie/>.
- Zhang, K., Chen, K., and Xue, W. (2011). Kangaroo: Reliable execution of scientific applications with dag programming model. In *Parallel Processing Workshops (ICPPW), 2011 40th International Conference on*, pages 327–334. IEEE.