

APPARATUS AND METHOD FOR SCHEDULING DISTRIBUTED WORKFLOW TASKS

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims priority to U.S. Patent Application No. 14/506,500 filed October
5 3, 2014, the contents of which are incorporated herein by reference.

FIELD OF THE INVENTION

This invention relates generally to distributed data processing in computer networks.
More particularly, this invention relates to scheduling distributed workflow tasks in computer
networks.

BACKGROUND OF THE INVENTION

Map Reduce is a programming model and an associated implementation for processing
and generating large data sets with a parallel, distributed algorithm executing on a cluster of
computers. Map Reduce utilizes distributed servers to run various tasks in parallel, while
managing all communications and data transfers between the various parts of the system. This
15 provides for redundancy and fault tolerance. The Apache Hadoop platform is an open-source
software framework that implements Map Reduce. Data stored in a distributed file system are
assumed to be Hadoop Distributed File System (HDFS) or its derivatives.

In-memory data processing is fast. Apache Spark provides primitives for in-memory
cluster computing that allows user programs to load data into a cluster's memory and query it
20 repeatedly. Apache Spark is not tied to the Map Reduce paradigm and has far greater
performance in certain circumstances.

It would be desirable to leverage the strengths of in-memory data processing and Map
Reduce while executing complex analytics workflows.

SUMMARY OF THE INVENTION

25 According to an aspect, a server has a processor and a memory storing instructions
executed by the processor to access scheduling tools including an entity workload profile with a
work flow tasks schedule and work flow task dependencies. Processed data associated with a
work flow task within the entity workload profile is identified. The work flow task dependencies
are

analyzed to alter the work flow tasks schedule to prioritize work flow tasks that rely upon the processed data. A work flow task running on a parallel, distributed Map Reduce cluster of computers to a work flow task running on an in-memory cluster of computers is selectively transitioned, wherein the in-memory cluster of computers has data loaded into cluster memory that is repeatedly queried.

BRIEF DESCRIPTION OF THE FIGURES

The invention is more fully appreciated in connection with the following detailed description taken in conjunction with the accompanying drawings, in which:

FIGURE 1 illustrates a system to implement operations of the invention.

FIGURE 2 illustrates generally processing operations associated with an embodiment of the invention.

FIGURE 3 illustrates more particular processing operations associated with an embodiment of the invention.

FIGURE 4 illustrates DAG processing associated with an embodiment of the invention.

FIGURE 5 illustrates task prioritization based upon processed data.

Like reference numerals refer to corresponding parts throughout the several views of the drawings.

DETAILED DESCRIPTION OF THE INVENTION

Figure 1 illustrates a system 100 configured to implement operations of the invention.

The system includes a master node or coordinator node 102 and a set of slave or worker nodes 104_1 through 104_N used to implement a distributed computing environment. The computers are connected via a network 106, which may be any combination of wired and wireless networks. A client computer 107 is also connected to the network 106.

The master node 102 includes standard components, such as a central processing unit 110 and input/output devices 112 connected via a bus 114. The input/output devices 112 may include a keyboard, mouse, display and the like. A network interface circuit 116 is also connected to the bus 114 to provide connectivity to network 106. A memory 120 is also connected to the bus 114. The memory 120 stores a master module 122, which includes instructions executed by the central processing unit 110 to implement operations of the invention.

The slave computer 104_1 also includes standard components, such as a central processing unit 130, input/output devices 132, a bus 134 and a network interface circuit 136. A memory 140 is connected to bus 134. The memory 140 stores a slave module 142, which

What is claimed is:

1. A server, comprising:
a processor; and
a memory storing instructions executed by the processor to:
 - 5 access scheduling tools including an entity workload profile with a work flow tasks schedule and work flow task dependencies,
identify processed data associated with a work flow task within the entity workload profile,
analyze the work flow task dependencies to alter the work flow tasks schedule
 - 10 to prioritize work flow tasks that rely upon the processed data, and
selectively transition a work flow task running on a parallel, distributed Map Reduce cluster of computers to a work flow task running on an in-memory cluster of computers, wherein the in-memory cluster of computers has data loaded into cluster memory that is repeatedly queried.
- 15 2. The server of claim 1 wherein the scheduling tools include a data profile store with individual data profiles for individual data sources.
3. The server of claim 2 wherein the individual data profiles include column level and data set wide statistics.
4. The server of claim 2 wherein the individual data profiles include total record counts,
20 minimum, maximum and mean values for numeric and date columns and cardinality estimations indicating the number of unique values in a column.
5. The server of claim 2 wherein the individual data profiles include frequency estimations for the most frequent values.

6. The server of claim 2 wherein the individual data profiles include a dataset footprint estimate.
7. The server of claim 2 further comprising instructions executed by the processor to estimate data growth and decay patterns for a work flow task.
- 5 8. The server of claim 1 wherein the scheduling tools include a preview engine to process a subset of data associated with a work flow task to develop analytics on the execution of the work flow task in its entirety.
9. The server of claim 1 wherein the scheduling tools include information on cluster resource availability.
- 10 10. The server of claim 1 wherein the scheduling tools include an operator composition analyzer to provide fine grain analysis of operators associated with work flow tasks to optimize the scheduling and execution of the operators.
11. The server of claim 1 wherein the scheduling tools include historical task execution profiles.