

(21) Application No: **1417283.7**
 (22) Date of Filing: **30.09.2014**
 (30) Priority Data:
 (31) **14043753** (32) **01.10.2013** (33) **US**

(51) INT CL:
G06F 17/30 (2006.01)
 (56) Documents Cited:
None
 (58) Field of Search:
 Other: **No search performed: Section 17(5)(b)**

(71) Applicant(s):
Cloudera Inc
Bldg 2, 1001 Page Mill Road, Palo Alto 94304,
California, United States of America
 (72) Inventor(s):
Marcel Kornacker
Justin Erickson
Nong Li
Lenni Kuff
Henry Noel Robinson
Alan Choi
Alex Behm
 (74) Agent and/or Address for Service:
Mewburn Ellis LLP
33 Gutter Lane, LONDON, EC2V 8AS, United Kingdom

(54) Title of the Invention: **Background format optimization for enhanced SQL-like queries in hadoop**
 Abstract Title: **Background format optimization for enhanced sql-like queries in hadoop**

(57) A format conversion engine for Apache Hadoop that converts data from its original format to a database-like target format, such as the columnar Parquet format, at certain time points for use by a low latency (LL) query engine 114 a-c, comprises a daemon 118a-c that is installed on each data node 120a-c in a Hadoop cluster. The daemon comprises a scheduler and a converter. The scheduler determines when to perform the format conversion and notifies the converter when the time comes. The converter converts data on the data node from its original format to a database-like format, for use by the low latency (LL) query engine. The converted data is saved on the data node together with the original, unconverted data, both available to the low latency (LL) query engine. The low latency (LL) query engine divides a query from a client into fragments for the target format which are then distributed among the data nodes for execution in parallel. The query fragments are executed on the transformed data to obtain intermediate results that are aggregated and returned to the client.

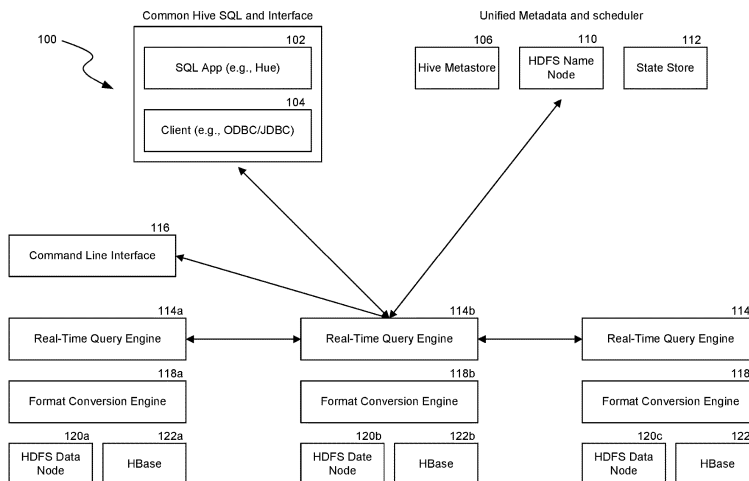


FIG. 1

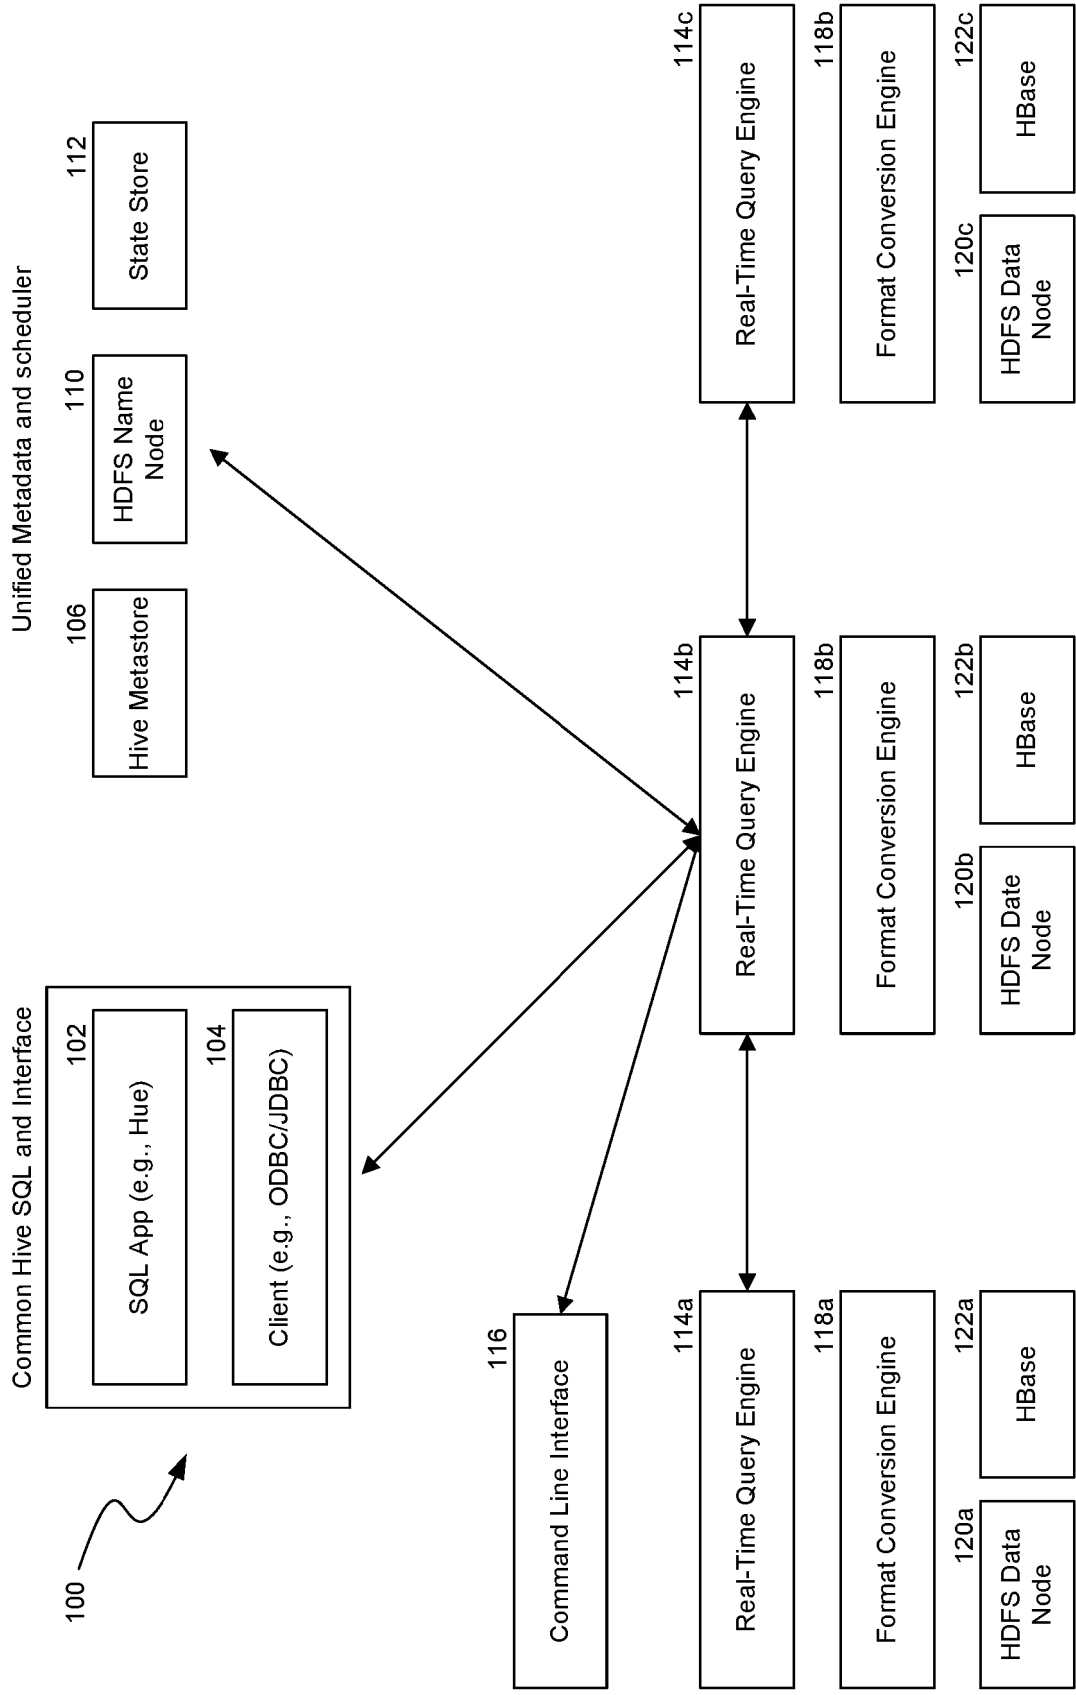


FIG. 1

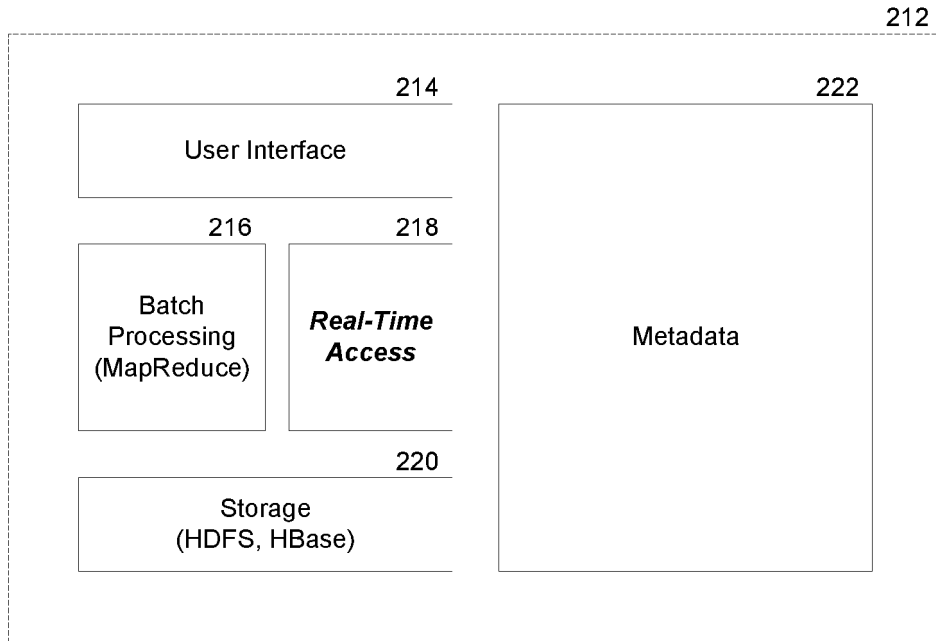


FIG. 2

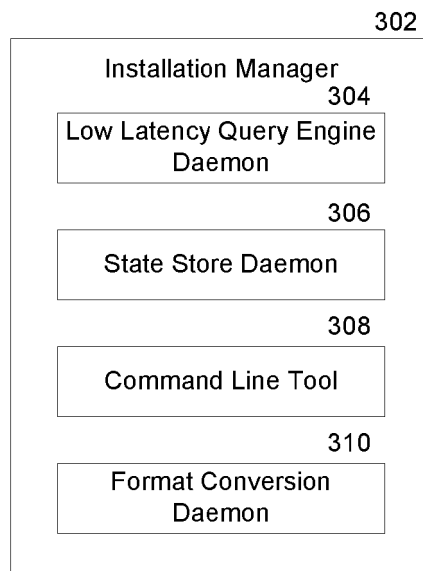


FIG. 3




FIG. 4

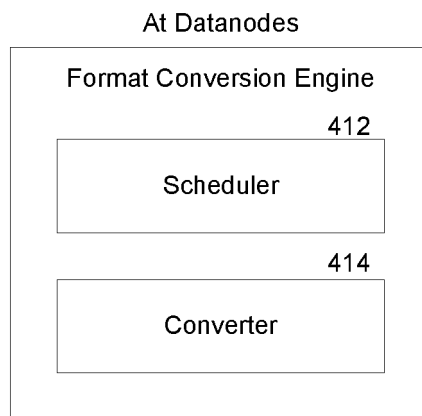
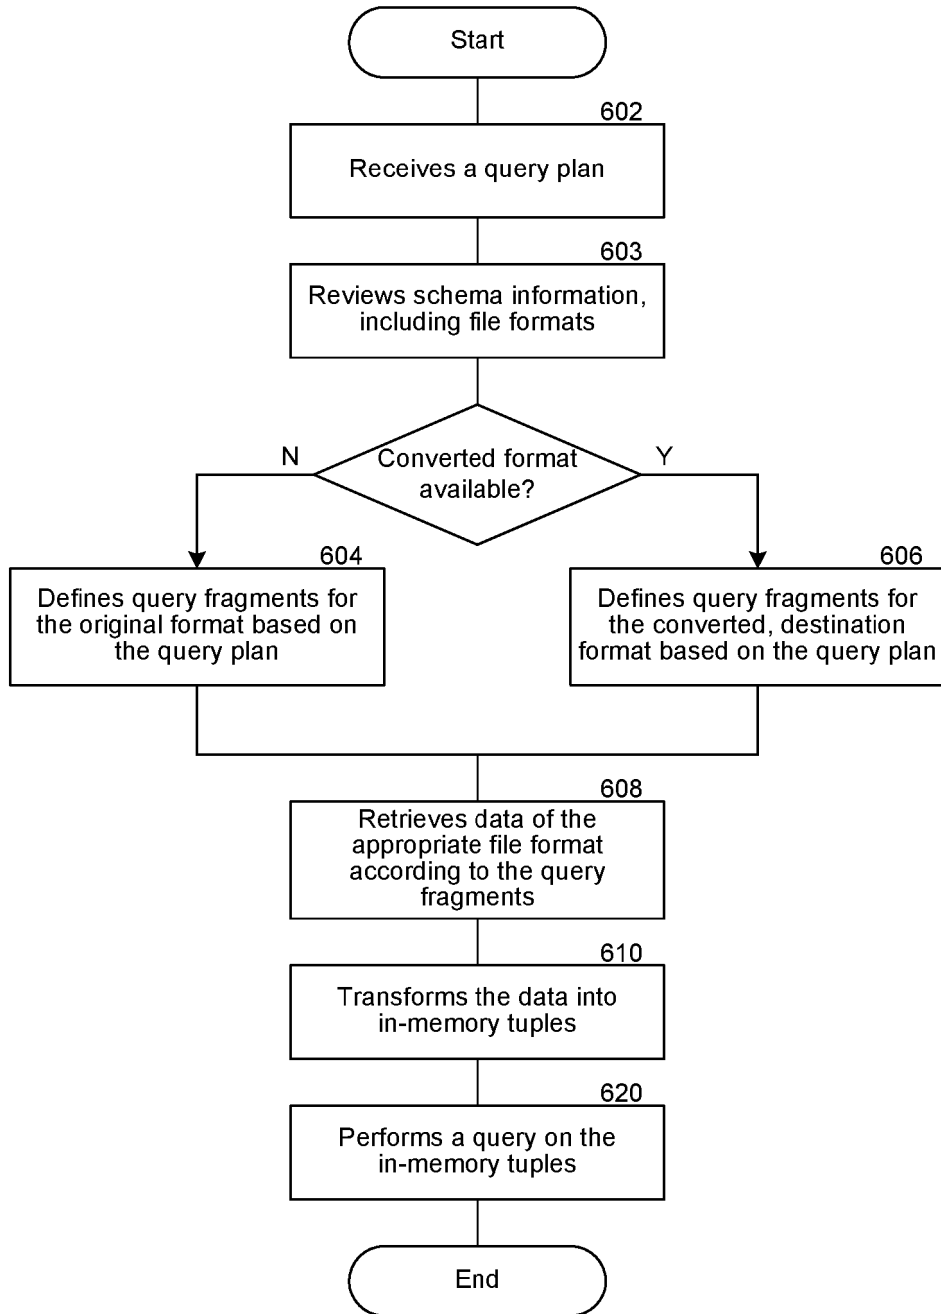


FIG. 5

**FIG. 6**

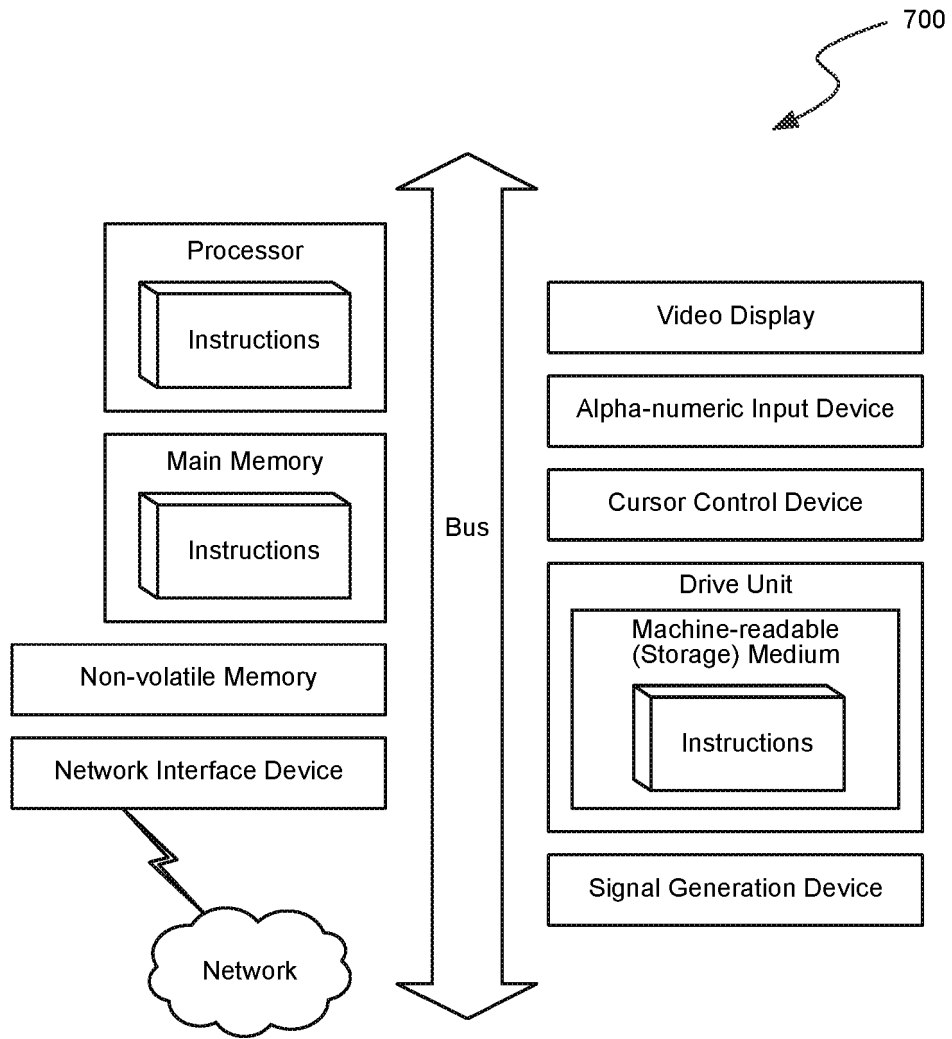


FIG. 7

BACKGROUND FORMAT OPTIMIZATION FOR ENHANCED SQL-LIKE QUERIES

IN HADOOP

BACKGROUND

[0001] Apache Hadoop project (hereinafter “Hadoop”) is an open-source software framework for developing software for reliable, scalable and distributed processing of large data sets across clusters of commodity machines. Hadoop includes a distributed file system, known as Hadoop Distributed File System (HDFS). HDFS links together the file systems on local nodes to form a unified file system that spans an entire Hadoop cluster. Hadoop also includes Hadoop YARN that provides a framework for job scheduling and cluster resource management that is utilized by a programming framework known as MapReduce. Hadoop is also supplemented by other Apache projects including Apache Hive (hereinafter “Hive”) and Apache HBase (hereinafter “HBase”). Hive is a data warehouse infrastructure that provides data summarization and ad hoc querying. HBase is a scalable, distributed NoSQL (No Structured Query Language) database or data store that supports structured data storage for large tables.

[0002] Hadoop currently does not support a relational database management system (RDBMS). For a relational database, a schema – the organization of data into tables having specific columns together with a set of integrity constraints ensuring compatibility between the columns of the tables – can be defined. A typical RDBMS implements a schema-on-write model, where a schema is enforced on data as the data is written into the database. Specifically, the data is reorganized and filtered using the integrity constraints before the data is stored into the database. A schema-on-write model works well for answering known questions. If a previously unknown question needs to be answered, new data may need to be captured. However, the RDBMS cannot accept new data that does not match the schema. To accommodate the new data, typically old data needs to be deleted from the database, the schema needs to be modified, and new data needs to be parsed and loaded into the database. In addition, data architects typically need to ensure that all the systems connected to the RDBMS work with the updated schema. This process of accommodating the new data can take a long time. Until then, the new data cannot be captured to answer the previously unknown question.

[0003] On the other hand, Hadoop currently follows a schema-on-read model, where a schema is not enforced on data until the data is read from the database. In this case, a schema generally specifies an organization of data in terms of file formats. As a result, the processing of the data can be decoupled from the storage of the data. Specifically, the underlying storage system in Hadoop can take files in their original format (e.g., tab-delimited text files, CSV, XML, JSON, images, etc.), while allowing an associated schema to be designed later and stored separately. In response to a query, the stored data is then transformed in-memory according to the separately stored schema. By virtue of the schema-on-read model, input data can be quickly updated in the database, which encourages users to experiment with different schemas.

[0004] The schema-on-read model and the schema-on-write model have their distinct merits. It would be useful for Hadoop to offer users the flexibility of using either one or a combination of them depending on the specific needs and requirements.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] Figure 1 contains a diagram illustrating an example environment in which a low latency query engine and a format conversion engine may be deployed.

[0006] Figure 2 contains a block diagram illustrating example components of a unified platform supporting batch-oriented and real-time, ad hoc queries.

[0007] Figure 3 contains a block diagram illustrating example components of an installation manager.

[0008] Figure 4 contains a block diagram illustrating example components of a low latency (LL) query engine daemon installed on each data node in a Hadoop cluster.

[0009] Figure 5 contains a block diagram illustrating example components of a format conversion engine daemon installed on each data node in a Hadoop cluster.

[0010] Figure 6 contains a flow diagram illustrating example operation of a query execution engine before it executes a collection of query fragments.

[0011] Figure 7 depicts a diagrammatic representation of a machine in the example form of a computer system within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed.

DETAILED DESCRIPTION

[0012] The following description and drawings are illustrative and are not to be construed as limiting. Numerous specific details are described to provide a thorough understanding of the disclosure. However, in certain instances, well-known or conventional details are not described in order to avoid obscuring the description. References to one or an embodiment in the present disclosure can be, but not necessarily are, references to the same embodiment; and, such references mean at least one of the embodiments.

[0013] Reference in this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the disclosure. The appearances of the phrase “in one embodiment” in various places in the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Moreover, various features are described which may be exhibited by some embodiments and not by others. Similarly, various requirements are described which may be requirements for some embodiments but no other embodiments.

[0014] The terms used in this specification generally have their ordinary meanings in the art, within the context of the disclosure, and in the specific context where each term is used. Certain terms that are used to describe the disclosure are discussed below, or elsewhere in the specification, to provide additional guidance to the practitioner regarding the description of the disclosure. For convenience, certain terms may be highlighted, for example using italics and/or quotation marks. The use of highlighting has no influence on the scope and meaning of a term; the scope and meaning of a term is the same, in the same context, whether or not it is highlighted. It will be appreciated that same thing can be said in more than one way.

[0015] Consequently, alternative language and synonyms may be used for any one or more of the terms discussed herein, nor is any special significance to be placed upon whether or not a term is elaborated or discussed herein. Synonyms for certain terms are provided. A recital of one or more synonyms does not exclude the use of other synonyms. The use of examples anywhere in this specification including examples of any terms discussed herein is illustrative only, and is not

intended to further limit the scope and meaning of the disclosure or of any exemplified term. Likewise, the disclosure is not limited to various embodiments given in this specification.

[0016] Without intent to further limit the scope of the disclosure, examples of instruments, apparatus, methods and their related results according to the embodiments of the present disclosure are given below. Note that titles or subtitles may be used in the examples for convenience of a reader, which in no way should limit the scope of the disclosure. Unless otherwise defined, all technical and scientific terms used herein have the same meaning as commonly understood by one of ordinary skill in the art to which this disclosure pertains. In the case of conflict, the present document, including definitions will control.

[0017] Embodiments of the present disclosure include a format conversion engine for Hadoop. Embodiments of the present disclosure also include systems and methods for performing format conversion, in real time or near real time, on data stored in Hadoop.

[0018] In one embodiment, the format conversion engine provides a mechanism enabling fast searches by making data ready in an easily queryable format. In another embodiment, it provides users with the flexibility of being able to update data quickly and to work with stabilized data efficiently.

[0019] Figure 1 contains a diagram illustrating an example environment 100 in which a low latency (LL) query engine and a format conversion engine may be deployed. The environment 100 includes a plurality of data nodes 120a-c that comprise a Hadoop cluster. Some of the data nodes 120a-c may run just HDFS, while others may run HBase region servers 122a-c.

[0020] The environment 100 includes a client 104 such as Java Database Connectivity (JDBC) client, Open Database Connectivity (ODBC) client, and the like that provides API and other tools for connecting and/or accessing a Hadoop cluster. SQL applications 102 such as Hue, provide a user interface for Hadoop to run queries or jobs, browse the HDFS, create workflows and the like. The environment 100 also includes a command line interface 116 for issuing queries. In one embodiment, the client 104, the SQL application 102 and the command line interface 116, each or together may be commonly referred to as a client.

[0021] A low latency (LL) query engine daemon 114 a-c runs on each of the data nodes. A low latency (LL) query engine daemon is a long running process that coordinates and executes queries. Each low latency (LL) query engine daemon 114a-c can receive, plan and coordinate queries received via the client's 102/104. For example, a low latency (LL) query engine daemon can divide a query into fragments, which are distributed among remote nodes running additional low latency (LL) query engine daemons for execution in parallel. The queries are executed directly on the HDFS (e.g., 120a-c) and/or HBase (e.g., 122a-c).

[0022] A format conversion engine daemon 118a-c also runs on each of the data nodes. The format conversion engine daemon 118a-c is a long running process that converts data from its original format to a condensed format that is conducive to relational database processing, such as the columnar format Parquet. The conversion can be performed at one or more time points. The converted data is saved on the data node together with the original, unconverted data, both available to the low latency (LL) query engine.

[0023] The environment 100 further includes unified metadata components such as a Hive metastore 106, an HDFS name node 110 and/or a state store 112. The Hive metastore 106 includes information about the data available to the various engines within the environment 100. Specifically, the Hive metastore 106 includes the schemas for the data stored on the data nodes 120a-c. The HDFS name node (NN) 110 includes the details of the distribution of files across the data nodes 120a-c to optimize local reads. In one implementation, the name node 110 may include information concerning disk volumes the files sit on, on an individual node.

[0024] The state store 112 is a global system repository which runs on a single node in the cluster. The state store 112 in one implementation can be used as a name service. All low latency (LL) query engine daemons, at start up, can register with the state store to be a member and get existing membership information specifying all the low latency (LL) query engine daemons that are running on the cluster. The state store 112, in a further implementation, can be used to provide metadata for running queries. The state store 112 can cache metadata and distribute the metadata to the low latency (LL) query engine daemons at start up or another time. When the state store fails, the rest of the system may continue to operate based on last information received from the state store. In a further implementation, the state store can store and distribute other system

information such as load information, diagnostics information, and the like that may be used to improve the functioning and/or performance of the Hadoop cluster.

[0025] Figure 2 contains a block diagram illustrating example components of a unified Hadoop platform 212 supporting batch-oriented and real-time, ad hoc queries. The unified Hadoop platform 212 supports distributed processing and distributed storage. The unified Hadoop platform 212 includes a user interface 214, storage 220 and metadata 222 components. The user interface 214 includes Hive interfaces such as ODBC driver, JDBC driver, Hue Beeswax, and the like. The user interface 214 also includes SQL support. Via the user interface 214, queries can be issued, data can be read from or written to storage 220, etc. The storage 220 includes HDFS and/or HBase storage. The HDFS may support various file formats, including but not limited to: text file, sequence file, RC file, Avro, and the like. Various compression codecs including snappy, gzip, deflate, bzip, and the like may also be supported. The metadata 222 may include, for example, information about tables, their partitions, columns, types, table/block locations, and the like. The metadata 222 may leverage existing Hive metastore, which includes a mapping of HBase tables, predicates on row key columns mapped into start/stop row, predicates on other columns mapped into single column value filters, and the like.

[0026] Existing Hadoop platform uses a batch-oriented query engine (i.e., MapReduce) for batch processing 216 of Hadoop data. The batch processing capability of MapReduce is complemented by a real-time access component 218 in the unified Hadoop platform 212. The real-time access component 218 allows real-time, ad hoc SQL queries to be performed directly on the unified storage 220 via a distributed low latency (LL) query engine that is optimized for low-latency. The real-time access component 218 can thus support both queries and analytics on big data.

[0027] Figure 3 contains a block diagram illustrating example components of an installation manager 302 for installing components of the various engines in a Hadoop cluster to provide interactive, real-time SQL queries directly on a unified storage layer. The installation manager 302 can automatically install, configure, manage and monitor the various engines. Alternately, the engines may be installed manually. The installation manager 302 installs four binaries including a low latency (LL) query engine daemon 304, a state store daemon 306, a low latency (LL) query engine shell 308 and a format conversion engine daemon 310. As described above, the low latency

(LL) query engine daemon 304 is a service or process that plans and executes queries against HDFS and/or HBase data. It is installed on each data node in the cluster. The format conversion engine daemon is a service or process that converts data from its original format to a condensed format. It is also installed on each data node in the cluster. The state store daemon 306 is a name service that tracks the location and status of all the low latency (LL) query engine daemons in the cluster. The state store daemon 306 can also be a metadata store for providing metadata and/or other diagnostic information in some implementations. The low latency (LL) query engine shell 308 is a command line interface for issuing queries to a low latency (LL) query engine daemon, and is installed on a client.

[0028] Figure 4 contains a block diagram illustrating example components of a low latency (LL) query engine daemon installed on each data node in a Hadoop cluster. A low latency (LL) query engine daemon includes a query planner 316, a query coordinator 318 and a query execution engine 320 in one embodiment.

[0029] The query planner 316 turns query requests from clients into collections of plan fragments based on stored schemas, and provides the plan fragments to the query coordinator 318. The query planner 316 may constitute the front end of the low latency (LL) query engine daemon written in Java or another suitable language to facilitate interaction with the rest of the Hadoop environment, such as the Hive metastore, the state store, APIs, and the like. The query planner 316 can use various operators such as Scan, HashJoin, HashAggregation, Union, TopN, Exchange, and the like to construct a query plan. Each operator can either materialize or generate data or combine data in some way. In one implementation, for example, the query planner can create a lefty plan or tree of one or more operators (e.g., manually or using an optimizer). The scan operator allows a plan to be broken up along scan lines or boundaries. Specialized scan nodes may be present for different storage managers. For example, there may be an HDFS scan node and an HBase scan node, each of which can internally employ different processes for different file formats. Some plans combine data for hash aggregation which can fill up a hash table and then output the aggregate results. A union operator can merge the output from different plan fragments. A TopN operator can be the equivalent of order by with a limit. The exchange operator can handle the data exchange between two plan fragments running on two different nodes.

[0030] The query coordinator 318 initiates execution of the plan fragments across all the low latency (LL) query engine daemons that are involved in the query. The query coordinator 318 uses the membership information from the state store and/or the location information for the data blocks from HDFS the Name Node to determine or identify the low latency (LL) query engine daemons for executing query plan fragments. In one implementation, the query coordinator 318 can also apply any predicates from the query to narrow down to the set of files and blocks against which the plan fragments should be run. The query coordinator 318 can also perform the final aggregation or merge of data from the low latency (LL) query engine daemons on remote data nodes. In one implementation, the low latency (LL) query engine daemons may pre-aggregate some of the data, thereby distributing the aggregation across the data nodes and speeding up the query processing.

[0031] The query execution engine 320 executes the planned query fragments locally on the HDFS and HBase. For example, it runs the scan and/or any other query operators. The query execution engine 320 is written in C++, but may also be written in any other suitable language, such as Java. The query execution engine 320 is an execution engine that is separate from MapReduce. While the query execution engine 320 accesses the infrastructure that provides the data (e.g., HDFS and HBase), it does not utilize any of the infrastructures that support map reductions, such as job trackers and task trackers.

[0032] In one embodiment, initially, data comes in and is stored in their original format on the HDFS data nodes. One or more associated schemas comprising information on file formats in which data is stored, which can be created by a user or an administrator, are saved separately in the Hive metastore 106, at the same time as the data is stored or at a later time. In one embodiment, after a query is submitted, a query execution engine 320 on a data node which is to execute certain planned query fragments locally first transforms the files on the data node according to the schemas. Specifically, the query execution engine 320 reads a schema, which contains information on row and column endings, for example, for the files from the Hive metastore. It then reads the files from the data node, parses them in accordance with the file formats specified in the schema, and transforms the parsed data into a series of in-memory tuples according to further information in the schema. At that time, the query execution engine 320 is ready to execute the planned query fragments locally against the transformation result.

[0033] In one embodiment, the query execution engine 320 can include a low level virtual machine (LLVM) 322, an optimizer, or other compiler infrastructure, for run-time code generation in order to transform interpretive code into a format that can be efficiently executed by the central processing unit (CPU). A typical RDBMS, for instance, has interpretive code for evaluating expressions to extract data from indices and the like. The query execution engine 320 handles this issue by using low level virtual machines (LLVMs) to more tightly couple code with hardware. For example, an expression where A equals B over A+B equals C in a query can be evaluated by making three function calls. Instead of making the three function calls, an LLVM uses the operations that the CPU provides in order to evaluate the expression and achieve speed gains.

[0034] In a further embodiment, the low latency (LL) query engine daemon can also use special CPU instructions, in order to, for example, perform text processing and/or other resource intensive processes. By way of another example, hash value computations may be performed using a special Cyclic Redundancy Check (CRC32) instruction to achieve speed gains.

[0035] In one embodiment, the low latency (LL) query engine provides the advantage of low latency which allows users to query large volumes of data and obtain answers at much faster speed than possible using the existing batch processing framework of Hive and MapReduce. In a further embodiment, the real-time query engine provides flexibility in applying schemas used to search for hidden insights in large volumes of data.

[0036] It takes different amounts of time for a query execution engine to parse and transform data in different file formats. In general, the amount of time decreases when the file format is more conducive to relational database processing in response to SQL-like queries. Therefore, a format conversion engine converts data to such a file format in the background to increase the efficiency of query processing at runtime. Figure 5 contains a block diagram illustrating example components of a format conversion engine daemon installed on each data node in a Hadoop cluster. In one embodiment, the format conversion engine daemon includes a scheduler 412 and a converter 414. The scheduler 412 determines when to perform the format conversion based on input by an administrator or a user, and notifies the converter when the time comes. In one example, the scheduler 412 uses a timer for performing the format conversion periodically or at certain points in the future. The certain point in the future could be measured from the occurrence

of an event, such as the creation, initial update, or last update of the data. In other examples, the conversion is performed when the data has been updated, searched, searched with the same queries, and so on, for a certain number of times. Accordingly, the scheduler 412 keeps a counter of the total number of updates, of all queries, of specific queries, of distinct queries, and so on, so that the format conversion can be performed when the criteria involving these numbers are met. In further examples, the status of resource utilization on the data node is taken into consideration in scheduling the format conversion.

[0037] In one embodiment, the scheduler 412 maintains one schedule for each piece of data on the data node, for each original format, for each target format, for each pair of an original format and a target format, etc. In another embodiment, the scheduler 412 determines when to delete the conversion results from the data node, which can be similar to the determination of when to perform the format conversion, and notifies the converter 414 when the time comes. While a scheduler 412 of a format conversion engine daemon may work independently, it can also coordinate with the scheduler of another format conversion engine daemon to perform the format conversion in a systematic fashion across multiple data nodes or even the entire cluster.

[0038] The converter 414 performs the format conversion upon receiving a notification from the scheduler 412. In one embodiment, the converter 414 maintains a list of one or more target formats. It converts the data on the data node to one of the target formats based on input by an administrator a user, and saves the converted data on the data node along with the original data. For example, the converter 414 may read a file in the CSV format from the data node into memory, parse the file in accordance with the CSV format, convert it into a chosen Parquet format, and saves the file in the Parquet format on the data node together with the file in the CSV format. In one embodiment, the conversion may be fully automated between certain original formats and target formats, possibly based on specific schemas stored in the Hive metastore. For instance, every field in a CSV file can be automatically converted into a column in a Parquet file. The conversion may also be customized by an administrator or a user, who may decide to convert an input file into multiple output files in the same target format or different ones, each having select fields in the input file arranged in a specific order, for example. In another embodiment, the converter 414 also deletes certain conversion results upon receiving a notification from the scheduler 412.

[0039] As a target format is typically a condensed format that is conducive to relational database processing, having data ready in a target format speeds up processing of SQL-like queries. As the format conversion is performed at carefully selected time points in the background, it tends to minimize the use of resources and interference with other operations on the data nodes.

[0040] With the format conversion engine daemon, in one embodiment, after a query is submitted, a query planner would set up the plan fragments to indicate that converted data is available. The query execution engine on a data node then no longer needs to perform a complex transformation of the data on the data node. It can simply read the converted data from the data node, which would essentially be in a tuple form. The format conversion engine daemon therefore provides some benefits of the schema-on-write model by reducing the processing time when the data is used in query processing, without suffering some costs of the model, which requires a large processing time when the data is uploaded and updated.

[0041] Figure 6 contains a flow diagram illustrating example operations of query planning and execution in the presence of a format conversion engine. At step 602, a query planner receives a query. At step 603, the query planner reviews relevant schema information to identify the available file formats in which data is stored. If only data in an original format is available, at step 604, the query planner defines plan fragments for the original format. If data in a converted target format is also available, however, at step 606, the query planner defines plan fragments for the target format.

[0042] Upon receiving a collection of planned query fragments, at step 608, a query execution engine on a data node reads data in the appropriate file format from the data node. At step 610, the query execution engine transforms the data into a series of in-memory tuples according to the schema information. At step 612, the query execution engine executes the collection of planned query fragments using the in-memory tuples. By virtue of these features, a user is given the flexibility of being able to experiment with datasets having different structures without incurring much overhead in data upload and update while being able to extract specific insight from the datasets in an efficient manner.

[0043] Figure 7 shows a diagrammatic representation of a machine in the example form of a computer system within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed.

[0044] In the example of Figure 7, the computer system 700 includes a processor, memory, non-volatile memory, and an interface device. Various common components (e.g., cache memory) are omitted for illustrative simplicity. The computer system 700 is intended to illustrate a hardware device on which any of the components depicted in the example of Figure 1 (and any other components described in this specification) can be implemented. The computer system 700 can be of any applicable known or convenient type. The components of the computer system 700 can be coupled together via a bus or through some other known or convenient device.

[0045] The processor may be, for example, a conventional microprocessor such as an Intel Pentium microprocessor or Motorola power PC microprocessor. One of skill in the relevant art will recognize that the terms “machine-readable (storage) medium” or “computer-readable (storage) medium” include any type of device that is accessible by the processor.

[0046] The memory is coupled to the processor by, for example, a bus. The memory can include, by way of example but not limitation, random access memory (RAM), such as dynamic RAM (DRAM) and static RAM (SRAM). The memory can be local, remote, or distributed.

[0047] The bus also couples the processor to the non-volatile memory and drive unit. The non-volatile memory is often a magnetic floppy or hard disk, a magnetic-optical disk, an optical disk, a read-only memory (ROM), such as a CD-ROM, EPROM, or EEPROM, a magnetic or optical card, or another form of storage for large amounts of data. Some of this data is often written, by a direct memory access process, into memory during execution of software in the computer 800. The non-volatile storage can be local, remote, or distributed. The non-volatile memory is optional because systems can be created with all applicable data available in memory. A typical computer system will usually include at least a processor, memory, and a device (e.g., a bus) coupling the memory to the processor.

[0048] Software is typically stored in the non-volatile memory and/or the drive unit. Indeed, for large programs, it may not even be possible to store the entire program in the memory.

Nevertheless, it should be understood that for software to run, if necessary, it is moved to a computer readable location appropriate for processing, and for illustrative purposes, that location is referred to as the memory in this paper. Even when software is moved to the memory for execution, the processor will typically make use of hardware registers to store values associated with the software, and local cache that, ideally, serves to speed up execution. As used herein, a software program is assumed to be stored at any known or convenient location (from non-volatile storage to hardware registers) when the software program is referred to as “implemented in a computer-readable medium.” A processor is considered to be “configured to execute a program” when at least one value associated with the program is stored in a register readable by the processor.

[0049] The bus also couples the processor to the network interface device. The interface can include one or more of a modem or network interface. It will be appreciated that a modem or network interface can be considered to be part of the computer system. The interface can include an analog modem, isdn modem, cable modem, token ring interface, satellite transmission interface (e.g. “direct PC”), or other interfaces for coupling a computer system to other computer systems. The interface can include one or more input and/or output devices. The I/O devices can include, by way of example but not limitation, a keyboard, a mouse or other pointing device, disk drives, printers, a scanner, and other input and/or output devices, including a display device. The display device can include, by way of example but not limitation, a cathode ray tube (CRT), liquid crystal display (LCD), or some other applicable known or convenient display device. For simplicity, it is assumed that controllers of any devices not depicted in the example of Figure 8 reside in the interface.

[0050] In operation, the computer system 800 can be controlled by operating system software that includes a file management system, such as a disk operating system. One example of operating system software with associated file management system software is the family of operating systems known as Windows® from Microsoft Corporation of Redmond, Washington, and their associated file management systems. Another example of operating system software with its associated file management system software is the Linux operating system and its associated file management system. The file management system is typically stored in the non-volatile memory and/or drive unit and causes the processor to execute the various acts required by the operating

system to input and output data and to store data in the memory, including storing files on the non-volatile memory and/or drive unit.

[0051] Some portions of the detailed description may be presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0052] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as “processing” or “computing” or “calculating” or “determining” or “displaying” or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0053] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the methods of some embodiments. The required structure for a variety of these systems will appear from the description below. In addition, the techniques are not described with reference to any particular programming language, and various embodiments may thus be implemented using a variety of programming languages.

[0054] In alternative embodiments, the machine operates as a standalone device or may be connected (e.g., networked) to other machines. In a networked deployment, the machine may operate in the capacity of a server or a client machine in a client-server network environment, or as a peer machine in a peer-to-peer (or distributed) network environment.

[0055] The machine may be a server computer, a client computer, a personal computer (PC), a tablet PC, a laptop computer, a set-top box (STB), a personal digital assistant (PDA), a cellular telephone, an iPhone, a Blackberry, a processor, a telephone, a web appliance, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine.

[0056] While the machine-readable medium or machine-readable storage medium is shown in an exemplary embodiment to be a single medium, the term “machine-readable medium” and “machine-readable storage medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and servers) that store the one or more sets of instructions. The term “machine-readable medium” and “machine-readable storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by the machine and that cause the machine to perform any one or more of the methodologies of the presently disclosed technique and innovation.

[0057] In general, the routines executed to implement the embodiments of the disclosure, may be implemented as part of an operating system or a specific application, component, program, object, module or sequence of instructions referred to as “computer programs.” The computer programs typically comprise one or more instructions set at various times in various memory and storage devices in a computer, and that, when read and executed by one or more processing units or processors in a computer, cause the computer to perform operations to execute elements involving the various aspects of the disclosure.

[0058] Moreover, while embodiments have been described in the context of fully functioning computers and computer systems, those skilled in the art will appreciate that the various embodiments are capable of being distributed as a program product in a variety of forms, and that the disclosure applies equally regardless of the particular type of machine or computer-readable media used to actually effect the distribution.

[0059] Further examples of machine-readable storage media, machine-readable media, or computer-readable (storage) media include but are not limited to recordable type media such as volatile and non-volatile memory devices, floppy and other removable disks, hard disk drives, optical disks (e.g., Compact Disk Read-Only Memory (CD ROMS), Digital Versatile Disks, (DVDs), etc.), among others, and transmission type media such as digital and analog communication links.

[0060] Unless the context clearly requires otherwise, throughout the description and the claims, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of “including, but not limited to.” As used herein, the terms “connected,” “coupled,” or any variant thereof, means any connection or coupling, either direct or indirect, between two or more elements; the coupling of connection between the elements can be physical, logical, or a combination thereof. Additionally, the words “herein,” “above,” “below,” and words of similar import, when used in this application, shall refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word “or,” in reference to a list of two or more items, covers all of the following interpretations of the word: any of the items in the list, all of the items in the list, and any combination of the items in the list.

[0061] The above detailed description of embodiments of the disclosure is not intended to be exhaustive or to limit the teachings to the precise form disclosed above. While specific embodiments of, and examples for, the disclosure are described above for illustrative purposes, various equivalent modifications are possible within the scope of the disclosure, as those skilled in the relevant art will recognize. For example, while processes or blocks are presented in a given order, alternative embodiments may perform routines having steps, or employ systems having blocks, in a different order, and some processes or blocks may be deleted, moved, added, subdivided, combined, and/or modified to provide alternative or subcombinations. Each of these processes or blocks may be implemented in a variety of different ways. Also, while processes or blocks are at times shown as being performed in series, these processes or blocks may instead be performed in parallel, or may be performed at different times. Further any specific numbers noted herein are only examples: alternative implementations may employ differing values or ranges.

[0062] The teachings of the disclosure provided herein can be applied to other systems, not necessarily the system described above. The elements and acts of the various embodiments described above can be combined to provide further embodiments.

[0063] Any patents and applications and other references noted above, including any that may be listed in accompanying filing papers, are incorporated herein by reference. Aspects of the disclosure can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above to provide yet further embodiments of the disclosure.

[0064] These and other changes can be made to the disclosure in light of the above Detailed Description. While the above description describes certain embodiments of the disclosure, and describes the best mode contemplated, no matter how detailed the above appears in text, the teachings can be practiced in many ways. Details of the system may vary considerably in its implementation details, while still being encompassed by the subject matter disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the disclosure should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features, or aspects of the disclosure with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the disclosure to the specific embodiments disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the disclosure encompasses not only the disclosed embodiments, but also all equivalent ways of practicing or implementing the disclosure under the claims.

[0065] While certain aspects of the disclosure are presented below in certain claim forms, the inventors contemplate the various aspects of the disclosure in any number of claim forms. For example, while only one aspect of the disclosure is recited as a means-plus-function claim under 35 U.S.C. §112, ¶13, other aspects may likewise be embodied as a means-plus-function claim, or in other forms, such as being embodied in a computer-readable medium. (Any claims intended to be treated under 35 U.S.C. §112, ¶13 will begin with the words “means for”.) Accordingly, the applicant reserves the right to add additional claims after filing the application to pursue such additional claim forms for other aspects of the disclosure.

CLAIMS

What is claimed is:

1. A system for performing queries on stored data in a distributed computing cluster of a plurality of data nodes, comprising:
 - a query engine for each data node, having:
 - a query planner that parses a query from a client to create query fragments based on a schema specifying one or more formats in which data is stored on the data nodes,
 - wherein, when data in a target format is stored, the query fragments are created for the target format, and when data in the target format is not stored, the query fragments are created for another format;
 - a query coordinator that distributes the query fragments among the plurality of data nodes; and
 - a query execution engine comprising:
 - a transformation module that transforms the data in the format for which the query fragments are created based on the schema; and
 - an execution module that executes the query fragments on the transformed data to obtain intermediate results that are aggregated and returned to the client.
2. The system of claim 1, wherein the distributed computing cluster is a Hadoop cluster.
3. The system of claim 1, wherein the target format is a columnar format.
4. The system of claim 1, wherein the target format is optimized for relational database processing.
5. A method of data processing for query execution, comprising the steps of:
 - storing initial data in an original format;
 - converting the initial data to be in a target format that is optimized for relational database processing according to a predetermined schedule; and
 - storing the converted data together with the initial data.

6. The method of claim 5, further comprising the steps of:
receiving a set of one or more query fragments after the converted data is stored;
transforming the converted data in response to the receipt; and
executing the set of one or more query fragments on the transformed data.
7. The method of claim 5, wherein the predetermined schedule is periodic.
8. The method of claim 6, wherein the predetermined schedule is based on a number of sets of query fragments that have been received.
9. The method of claim 5, further comprising the step of
storing new data in the original format to replace the initial data,
wherein the predetermined schedule is based on a number of times stored data in the original format has been replaced.
10. The method of claim 5, wherein the target format is a columnar format.
11. The method of claim 5, further comprising the step of deleting the converted data according to a schedule.
12. The method of claim 6,
wherein the initial data and the converted data are stored on a data node in a distributed computing cluster, and
wherein the set of one or more query fragments are executed on the data node.
13. A system for data processing for query execution, comprising:
a first storing unit which stores initial data in an original format;
a converting unit which converts the initial data to be in a target format that is optimized for relational database processing according to a predetermined schedule; and
a second storing unit which stores the converted data.

14. The system of claim 13, further comprising:
 - a receiving unit which receives a set of one or more query fragments after the converted data is stored;
 - a transforming unit which transforms the converted data in response to the receipt; and
 - an executing unit which executes the set of one or more query fragments on the transformed data.
15. The system of claim 13, wherein the target format is a columnar format.
16. The system of claim 13, wherein the system is a node in distributed computing cluster.
17. A machine-readable storage medium having stored thereon instructions which when executed by one or more processors perform a method, the method comprising the steps of:
 - storing initial data in an original format;
 - converting the initial data to be in a target format that is optimized for relational database processing according to a predetermined schedule; and
 - storing the converted data.
18. The machine-readable storage medium of claim 17, the method further comprising the steps of:
 - receiving a set of one or more query fragments after the converted data is stored;
 - transforming the converted data in response to the receipt; and
 - executing the set of one or more query fragments on the transformed data.
19. The machine-readable storage medium of claim 17, wherein the target format is a columnar format.