



- (51) **International Patent Classification:**
G06F 17/40 (2006.01) G06F 15/16 (2006.01)
G06F 3/06 (2006.01)
- (21) **International Application Number:**
PCT/US2016/022968
- (22) **International Filing Date:**
17 March 2016 (17.03.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/134,370 17 March 2015 (17.03.2015) US
- (71) **Applicant:** CLOUDERA, INC. [US/US]; 1001 Page Mill Road, Bldg. 2, Palo Alto, CA 94306 (US).
- (72) **Inventor:** LIPCON, Todd; 500 Beale St #205, San Francisco, CA 94105 (US).
- (74) **Agents:** GOSWAMI, Tathagata D. et al.; Perkins Coie LLP, P.O. Box 1247, Seattle, WA 98111-1247 (US).

- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: COMPACTION POLICY

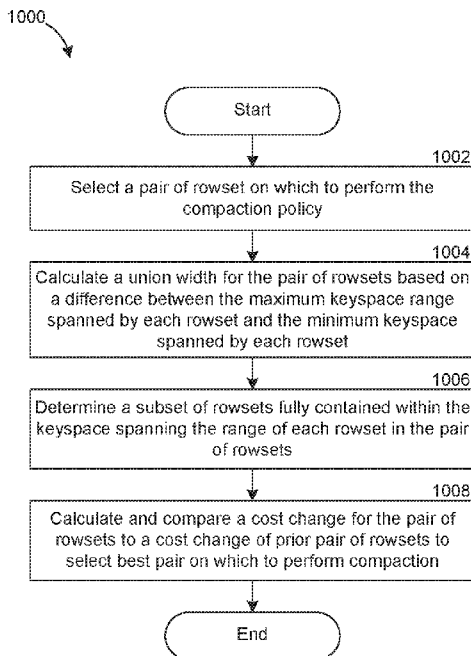



FIG. 10

(57) **Abstract:** A compaction policy imposing soft limits to optimize system efficiency is used to select various rowsets on which to perform compaction, each rowset storing keys within an interval called a keyspace. For example, the disclosed compaction policy results in a decrease in a height of the tablet, removes overlapping rowsets, and creates smaller sized rowsets. The compaction policy is based on the linear relationship shared between the keyspace height and the cost associated with performing an operation (e.g., an insert operation) in that keyspace. Accordingly, various factors determining which rowsets are to be compacted, how large the compacted rowsets are to be made, and when to perform the compaction, are considered within the disclosed compaction policy. Furthermore, a system and method for performing compaction on the selected datasets in a log-structured database is also provided.



Published:

— with international search report (Art. 21(3))

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

COMPACTION POLICY

COMPACTION POLICY

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Patent Application No. 62/134,370, filed March 17, 2015, which is incorporated herein by reference in its entirety.

BACKGROUND

[0002] The need for better data storage services and those that provide greater storage capacity have substantially increased in recent times. Furthermore, as a centralized approach to data storage becomes more prevalent, distributed databases such as those designed using cloud-based storage systems have become an industry standard.

[0003] In distributed, large-scale storage systems, improved indexing mechanisms are usually implemented to decrease latency (i.e., time to access data). When performing a read operation, for example, a storage system can look for queried data using an in-memory index mapped to various data nodes distributed across a network. Projects like Apache's HBase, Google's BigTable, etc. provide both the software and the software framework for reliable, scalable and distributed processing of large data sets in a network of computers, or clusters, communicating over the Internet. A particular file system, (e.g., Google's table or Apache's cluster), typically comprises a name node (master server) and a plurality of data nodes (tablet servers). In some instances, one or more clusters can also be referred to as a distributed database in the distributed file system (DFS). A DFS is typically managed by a service provider who deploys a unified file system in which a name node (running a file sub-system of the unified file system) manages a plurality of data nodes.

[0004] In order to efficiently read and update the files in each distributed database, or datastore, a name node includes various other software components which determine when certain thresholds have been reached. For example, the thresholds can be related to a time limit, a file size limit, and the like, for job scheduling and resource management within the distributed database. Additionally, the name node can determine when certain

failovers have occurred in data nodes order to re-direct data processing to other data nodes and avoid data loss.

[0005] For example, Apache's Zookeeper provides a centralized service for maintaining configuration information and naming, and also for providing distributed synchronization and group services. Apache's HBase provides a non-relational datastore, which is a scalable, distributed No-Structured-Query-Language (NoSQL) database that supports structured storage of large tables, similar to Google's BigTable. Generally, HBase includes a region server instance on each of the data nodes and depends on a ZooKeeper service running on the name node to coordinate the region servers. Each of the region servers manages data files underlying the large tables, look-ahead logs and other metadata in a data directory on the data node. Each of the distributed databases are also supplemented by additional projects which help provide programming framework for job scheduling and cluster resource management. Examples of additional projects can be Apache's MapReduce or Google's Chubby.

[0006] In order to maintain a table up-to-date, old data within the table needs to be removed and new data needs to be added quickly and efficiently. Furthermore, the data stored in the table needs to occupy the least amount of space possible in order to make room for the new data to be stored. Various techniques for updating and maintaining efficient tables are available. Such techniques are typically known as compaction.

[0007] Traditional compaction techniques, however, typically cause system resources to be wasted by performing compactions too often, or performing compactions on datasets which do not optimize tablet efficiency. This can substantially impact existing workloads as well as storage space for the DFS, which increases costs for the service provider. Thus, embodiments of the present disclosure facilitate an efficient approach for maximizing storage as well as minimizing the cost of I/O operations without causing interruptions.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Fig. 1 illustrates an exemplary architecture of a distributed filed system in which data compaction may be utilized.

[0009] Fig. 2 illustrates a block diagram including exemplary components of a compaction engine in the distributed file system of Fig. 1.

[0010] Fig. 3A illustrates an exemplary tablet of data stored in a distributed file system of Fig. 1.

[0011] Fig. 3B illustrates an exemplary RowSet within the table of Fig. 3A.

[0012] Fig. 4 illustrates an exemplary timing diagram for a process of performing a compaction on a tablet within a distributed file system.

[0013] Fig. 5 illustrates exemplary rowsets with overlapping key values.

[0014] Fig. 6 illustrates exemplary rowsets with bounded sizes.

[0015] Figs. 7-9 illustrate exemplary rowsets in a keyspace on which a compaction policy may be implemented.

[0016] Fig. 10 illustrates an exemplary method for selecting rowsets by a compaction policy.

[0017] Fig. 11 illustrates a screenshot of a compaction policy log demonstrating compaction.

[0018] Fig. 12 shows an exemplary computer system architecture for performing one or more of the methodologies discussed herein.

DETAILED DESCRIPTION

[0019] 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.

[0020] 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.

[0021] 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.

[0022] 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.

[0023] 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.

[0024] As used herein, a “server”, an “engine”, a “module”, a “unit” or the like may be a general purpose, dedicated or shared processor and/or, typically, firmware or software that is executed by the processor. Depending upon implementation-specific or other considerations, the server, the engine, the module or the unit can be centralized or its functionality distributed. The server, the engine, the module, the unit or the like can include general or special purpose hardware, firmware, or software embodied in a computer-readable (storage) medium for execution by the processor.

[0025] As used herein, a computer-readable medium or computer-readable storage medium is intended to include all mediums that are statutory (e.g., in the United States, under 35 U.S.C. 101), and to specifically exclude all mediums that are non-statutory in nature to the extent that the exclusion is necessary for a claim that includes the computer-readable (storage) medium to be valid. Known statutory computer-readable mediums include hardware (e.g., registers, random access memory (RAM), non-volatile (NV) storage, to name a few), but may or may not be limited to hardware.

[0026] In a log-structured database, instead of writing to a disk randomly, data is temporarily stored (e.g., buffered) in memory and a log is maintained to specify the disk location(s) where the data is to be written. This buffered data is flushed to the disk at a later time. Thus, there is a delay between the time when data comes in and when that data gets written onto the disk. This is typically because magnetic spinning disks are slow for random access operations. However, if data can be held temporarily in memory and then flushed out to disk, then sequential write operations can be performed on the disk. Because sequential write operations are faster than random write operations, this results in faster write operations.

[0027] As new data comes in, such data first enters a memory buffer referred to as MemRowSet, which is similar to an in-memory cache. The data in the buffer later spills to the disk, e.g., a datastore such as HBase, in smaller sorted files in a memory flush process. Thus, the data is typically sorted in memory, and then flushed on to the disk in smaller sorted files according to a memory flush process. The sorted files are typically

partitioned (horizontally) and stored in sorted string tables (SSTables) or Hfiles that are often referred to as tablets. These tablets are each portions of a much larger data table that is stored across a distributed file system (DFS), such as a Hadoop File System (HDFS) or BigTable.

[0028] Tablets can be subdivided into smaller units called RowSets. Some RowSets exist in memory only, termed MemRowSets, while others exist in a combination of disk and memory, termed DiskRowSets. However, it is likely that the primary key intervals of different RowSets may intersect or overlap. (Each rowset stores a collection of primary keys.) The previous MemRowSet is written to disk, and becomes one or more DiskRowSets. This flush process is fully concurrent: readers can continue to access the old MemRowSet while it is being flushed, and updates and deletes of rows in the flushing MemRowSet are carefully tracked and rolled forward into the on-disk data upon completion of the flush process.

[0029] In log-structured databases such as BigTable, Cassandra, Accumulo, and HBase, compaction techniques are used to organize the on-disk data into segments (or, equivalently referred to herein as chunks or rowsets or RowSets). The segments are ordered by a primary sort key. Accordingly, when servicing read workloads, these on-disk segments are accessed using the primary sort key to search for the data that is being read from the disk.

[0030] One idea behind compaction techniques is to facilitate a consistent or effective read performance. To maintain consistent read performance, the data segments, or the smaller sorted files are typically merged together into larger sorted segments. This improves the read performance at the time of reading/searching the data but, however, involves sequential I/O costs upfront in order to perform the merge. Thus, in order to service a read workload, the larger sorted segments are accessed instead of the plethora of smaller sorted segments that were written to the disk.

[0031] In compaction, smaller, individually sorted segments or rowsets are merged to larger sorted segments. For example, five or six of the smaller sorted segments can be merged together into a segment that is sorted using a primary sort key. Each segment or rowset can include a collection of rows and each row has a primary key. (In embodiments

where a segment includes more than one identical row, it is likely that the primary key appears in multiple rows within that segment.) Because the segments are sorted, algorithms such as binary search or B-tree index can be used to access, or read the sorted data, with a relatively smaller number of read operations.

[0032] In some embodiments, the contents (e.g., keys) of a segment (e.g., a rowset) are individually sorted but one segment is not related to other segments. Each rowset can have a plurality of rows, and each row has a primary sort key. For example, when searching for a primary sort key called "Todd," there could be one or more segment that includes this primary sort key. Each of these one or more (sorted) segment may not have a relationship with other segments. A logarithmic data access method such as binary search can be used to look up the data with the primary key "Todd" in each of the one or more individual (sorted) segments. Searching for the primary key "Todd" is of complexity $O(n)$, where n is the number of the one or more (sorted) segments that include this primary key.

[0033] Given the large amount of data stored on-disk, which can range in the petabytes across the DFS and megabytes on each tablet, compacting the smaller segments into larger segments results in a reduced read workload. However, servicing a read on a large sorted file, can cause latency since the length of a row (including multiple segments) can be significantly long and each segment in that row is to be sequentially accessed. Consequently, determining a middle ground to optimize the system efficiency thus becomes a balancing act of (a) selecting the appropriately-sized smaller sorted segments, (b) determining a size of the larger sorted segments, and (c) identifying the times when compaction is to be performed. For example, the disclosed compaction policy results in a decrease in a height of the tablet, removes overlapping rowsets, and creates smaller sized rowsets. Typically, rowsets that are selected for compaction are, on an average, those rowsets that overlap, are wide (e.g., span a wide keyrange) but are not too large in size.

[0034] Embodiments of the present disclosure provide systems and methods for performing compaction on log-structured databases. The systems and methods disclosed herein enable efficient data storage and retrieval for a distributed, e.g., cloud computing

environment. Specifically, taking into consideration the various factors relating to a data compression operation, can reduce the overall operating cost of the storage system. These factors facilitate data operations to increase data quality and availability to reduce I/O operations, with minimal impacts existing workloads.

[0035] Embodiments of the present disclosure also include systems and methods for performing a data flush into a log-structured distributed database as well as compaction methods for merging the existent data within the database. Embodiments of the present disclosure also provide methods for a compaction policy that selects how and which data is compacted on log-structured databases, such as within BigTable (Google), Cassandra (Apache/Dynamo), HBase (Apache/Hadoop), and similar non-relational distributed database, or DFS. These methods further define a compaction policy implemented to determine when to perform data compaction on a log-based database. The compaction policy includes a plurality of factors on which the compaction depends. For example, the disclosed methods can factor in the physical layout of the rowsets to improve the efficiency of subsequent operations. For example, in some scenarios, accessing multiple shorter rows of the smaller sorted files can be more efficient than accessing a significantly long row of these files. Also, compaction policies that result in minimum utilization of system resources are more desirable, thereby lowering I/O costs. In some exemplary embodiments, evenly distributing the system resources during the compaction process is desirable, such that system performance stays predictable and substantially constant. In other words, during the compaction process, it can be desired that the system performance shows no or relatively less spikes.

[0036] Specifically, embodiments of the present disclosure propose a new compaction policy for merging data files on-disk. The compaction policy pertains to merging rowsets in tablets. In exemplary embodiments, a compaction policy places a soft limit on the size of these files, and uses a cost-based analysis to select the optimal files that are to be merged. The compaction policy improves on previous methods by maximizing a user-directed measurement of performance while maintaining a given budget for I/O. Additionally, the compaction policy provides a more consistent performance by breaking all requests, e.g., the system workload, into size-bounded units.

[0037] The methods and system used to implement compaction in a distributed database, such as in a distributed file system (DFS) are further discussed in the following description, beginning with a general overview of a DFS with reference to FIG. 1

[0038] FIG. 1 illustrates an exemplary environment in which the disclosed compaction policy may be utilized to facilitate storage in a log-structure database system. In the illustrated environment, a non-relational distributed database (datastore) in a distributed file system (DFS) is provided. The datastore can also be referred to as a table or cluster, having a master node 100 and a plurality of slave nodes 108A, 108B, and 108C. The master node 100 may be implemented on one or more servers and may include file management software, such as a Zookeeper 114. A slave node 108A, for example, may include a data node 110A and a region server 112B instance. A data node 110A, for example, can comprise a plurality of tablets, which are portions of a large distributed table managed by the master node 100. Each tablet stores data which is tracked by the master node 100. Although not shown in Fig. 1, the master node 100 may also include an iterator, a garbage collector, and various other processes utilized within the datastore. The datastore shown in Fig. 1, may be deployed over a cloud-based platform via a cloud-based service provider, which may offer storage and/or other services remotely through a DFS. In general, the cloud-based service provider hosts a cloud-based platform which enables users to simultaneously access stored data remotely in real-time or near real-time. Users may additionally add, update, delete or otherwise modify the data stored in the cloud-based platform. The cloud-based service provider may implement an architecture suited for handling the large volume of user requests through the cloud-based platform. Accordingly, lowering I/O accesses and reducing storage on associated databases provides a more efficient system.

[0039] In general, the slave nodes 108 communicate with the master node 100 via a network. For example, an open network, such as the Internet, or a private network, such as an intranet and/or the extranet, or any combination or variation thereof. The network can be any collection of distinct networks operating wholly or partially in conjunction to provide connectivity to the DFS and may appear as one or more networks to the serviced systems and devices. In one embodiment, communications can be achieved by a secure

communications protocol, such as secure sockets layer (SSL), or transport layer security (TLS).

[0040] Fig. 2 shows a block diagram illustrating exemplary components of the compaction engine. In one embodiment, the compaction engine comprises a region manager module 202, a snapshot iterator module 204, and a compaction policy module 206. The snapshot iterator module 204 manages a snapshot for a table, including its creation and modification. It can operate in either an offline mode or an online mode.

[0041] Figs. 3A and 3B illustrate exemplary snapshots of a master table in which data is stored as different columns. For purposes of this description, DiskRowSets in a table are represented as "R' ", rowsets included within a DiskRowSet stored on a tablet are represented as "R", and rows included within a rowset are represented as "r". As shown, each DiskRowSet, e.g., represented as diskrowsets R'1, R'2, R'3, etc., can include multiple rowsets stored within a particular row and column in the master table. Thus, diskrowset R'1 includes rowsets R1, R2, R3, and R4. Diskrowset R'2 includes rows R5 and R6. Each rowset can include multiple rows. For example, rowset R1 includes rows r1, r2, r3, r4, and r5, as shown in Fig. 3B. (A DiskRowSet is a type of rowset, thus every DiskRowSet is also, in turn, a rowset.)

[0042] This method of encapsulation of data would allow, for example, DiskRowSet R'1 to be stored on a first tablet, and DiskRowSet R'2 to be stored on a second tablet. These tablets, can in turn, be managed by several region servers running on the data nodes of a cluster. For example, a region server 112A in tablet 108A of FIG. 1 can manage rowset R1 (i.e., row 304) and any addenda, and region server 112B in tablet 108B of FIG. 1 can manage files in diskrowset R'2 and any addenda, and so on. Each diskrowset, e.g., R'1, R'2, R'3 can be individually sorted, but they may not necessarily be sorted with respect to one another.

[0043] Similarly and as shown in Fig. 3B, each of C1 and C2 are representative of a column family under which numerous columns (c1, c2, etc.) can be associated. Accordingly, though the master table of the DFS may have any number of DiskRowSets (e.g., R'1, R'2) and columns (C1, C2), portions of the data included within that table can be distributed for handling across various networked computers, hence the name of a

distributed file system (DFS) and the distributed databases forming clusters within that DFS. Fig. 3B illustrates a table on a tablet, or data node, such as R'1 in Fig. 3A. The data files stored within a given datastore may therefore be stored in any one of these rows (r1), in a corresponding rowset (R1), and associated with a DiskRowSet row (R'1) in a table. The data file may also be associated with a particular column (c1) in a column family (C1).

[0044] Typically, in datastores such as HBase, changes to existing data are represented as addendums rather than modifying the existing data. Since the data is stored sequentially on a disk, when such data is accessed, the most recent version of the data is returned and the older version of the file is eventually discarded, or removed from disk during a garbage collection (GC) process.

[0045] Fig. 4 illustrates an exemplary compaction process, implemented in four stages on a particular rowset, e.g., such as on a data node in a cluster. The compaction of one or more rowsets can occur concurrently with an update to the rowsets. Compaction is also referred to herein as a merge of two or more rowsets on-disk, e.g., on a particular data node. In some embodiments, compaction can include a merge of two or more diskrowsets. In some embodiments, a compaction can include a merge of a diskrowset and a rowset.

[0046] Starting with the first stage 402 of compaction, a data node is initialized to a steady state. In the steady state, all updates, i.e., data added to the system or changes to already stored data, are applied only to a selected rowset or diskrowset. The selected rowset may also be referred to as a source rowset. For example, the source rowset may be DiskRowSet R'1, in Fig. 3A, which includes rowsets R1, R2, R3, and R4. In other embodiments, the source rowset may include, for example, rowset R1, or rowsets R2 and R4, or any one or combination thereof. In other embodiment, entire rowsets are compacted together, e.g., R1 and R2, rather than just rows r1, r2, r3, etc. within an individual rowset. For purposes of this discussion, and with reference to Fig 4, it is assumed that the rows within one rowset are compacted such that one input rowset (e.g., RowSet_{in}) is compacted to one output rowset (e.g., RowSet_{out}).

[0047] During a transition state 410, which occurs between the first stage 402 and the second stage 404, e.g., Phase I, the system may create a snapshot iterator to merge input

rowsets from an in-memory cache, or buffer, that store new data and updates to add to the source rowset. The input rowsets include rowsets of new data files that are to be stored on-disk at the corresponding data node, e.g., tablet. The snapshot iterator can also be utilized to save an associated snapshot taken during the transition state 410 which serves as a static copy of the current state of the source rowset prior to the merge. Accordingly, any reads on the source rowset occurring concurrently with the compaction process will reflect the same data until the merge is complete. This type of snapshot isolation is also commonly referred to in the art as a multiversion concurrency control (MVCC) snapshot of the source rowset at a current state. The MVCC snapshot can be saved on-disk as a backup to the source rowset state such that any data removed, changed and/or lost during a merge can be recovered and provided to the end-user, if necessary.

[0048] After the snapshot is saved during the transition state 410, the compaction process enters the second stage, which is the merge phase, "Phase I" 404 of the compaction process. In Phase I 404, the snapshot iterator created during the transition state 410, is utilized to create a new data set for an output rowset. The output rowset can include the combined data in the source rowset and in the input rowset (e.g., from the in-memory cache). In some embodiments, the output rowset can reflect updates or deletes which arrived prior to the compaction process entering Phase I 404. In some embodiments, updates or deletes occurring once Phase I 404 begins or during Phase II 406 are not reflected in the new data set of the output rowset.

[0049] In some embodiments, during Phase I 404, an additional integer column (also referred to herein as mapping column) is also created in order to maintain track of the rows in the input rowset. The integer column includes an integer entry for each row. Accordingly, if a row is garbage collected (GC) during the merge, the removal of those rows can be mapped to a corresponding row ID of the output rowset. For example, if the number of rows in the input rowset is "N" and no rows are compacted (e.g., deleted) in the output rowset, then the integer column will reflect "1, 2, 3, 4...N". On the other hand, if rows are compacted during the merge phase 404, then the integer column may reflect, e.g., "1, 2, x, 3, x, 4,...N", where each "x" is a value denoting a row deleted in the output rowset. Garbage collecting (GC) rows can provide additional on-disk space for new rowset

flushes from the in-memory cache and can help maintain the efficiency of the reads on-disk.

[0050] In some embodiments, if any mutations arrive during Phase I 404 those mutations are applied to the delta tracking structures, of the input rowset. Mutations, for example, include any changes to the data stored in the table. For example, adding data, deleting data, modifying data, re-arranging data, etc. are considered mutations. Since the merge in Phase I operates on the snapshot taken by the snapshot iterator, the mutations received during Phase I are not taken into account in the output rowset. In some embodiments, during a third stage "Phase II" 406 of the compaction process, these mutations are migrated into the output rowset.

[0051] In some embodiments, during the migration stage of Phase II 406, the mutations received during Phase I are handled, as well as any mutations arriving during Phase II. The mutations received during Phase I are simply migrated to the output rowset. This is accomplished by aggregating all the mutations not included in the snapshot iterator and merging the ordinal indexes with the mapping, or integer column which was created during Phase I. This merge effectively translates the old rowID to the new rowID. Accordingly, each mutation can be inserted into the output rowset's delta tracking structures.

[0052] To account for the mutation arriving in Phase II, the mutations are applied to both the input rowset and the output rowset. This replication is performed in order to ensure that no data loss occurs during the compaction process. For example, if a mutation is received for a GC'ed row as reflected in the integer column of the input rowset, then these mutations can be migrated into the subsequent row of the output rowset. The replication of the mutations can additionally help maintain data consistency and reliability within the system since any reads which occur during Phase II 406 of the compaction process are handled by the input rowset. Thus, any mutations arriving in Phase II 406 are still reflected in those reads.

[0053] In order to replicate the mutations received during Phase II 406 of the compaction process, the key lookup may be duplicated into the output rowset's key

column. Alternatively, in another embodiment, the integer column used for mapping the rows in Phase I may be utilized.

[0054] Once all of the mutations have been handled, the input rowset and the output rowset have logically identical data though the output rowset is compacted. Accordingly, Phase II ends and the process enters a fourth stage 408. In the fourth stage 408 the two logically identical rowsets (e.g., input rowset and output rowset) are atomically swapped. In other words, the rowset data is swapped without any interrupts occurring in the system and in a one-step operation. Once the output rowset is swapped with the input rowset, e.g., the output rowset effectively becomes the new input rowset and the input rowset becomes the new output rowset, any updates or other mutations arriving for the rowset are only applied to the new input rowset. The new output rowset (i.e., the original input rowset) and its associated mapping column may then be dropped, or discarded, providing additional disk space for new data to be stored. Then, the compaction process ends.

[0055] The compaction process discussed in Fig. 4 may also be applied to multiple rowsets concurrently. For example, during the transition state 410, instead of a snapshot being taken of just one rowset, a snapshot is taken of each rowset and a snapshot iterator is created. A merge iterator can also be created in order to merge all of the snapshots in ascending key order to a new data set corresponding to the multiple output rowsets. Similar to the compaction process discussed in Fig. 4, the merge iterator simultaneously creates an integer column for each rowset during the merge of each corresponding snapshot. The integer, or mapping column can include a rowID for each old rowset in a respective snapshot.

[0056] A compaction policy, or the process of selecting rowsets to apply the above compaction process in distributed database is now discussed with reference to Figs. 5-9. Though only one rowset or one keyspaces is discussed in certain embodiments, it is understood that the compaction policy may be applied to any number of keyspaces, combined or individually, as well as any number of rowsets, tablets, or tables.

[0057] Fig. 5 is an exemplary diagram of numerous rowsets of variable size in a plurality of keyspaces on, for example, a data node that includes a tablet 500. In embodiments of the present disclosure, tables and corresponding sub-tables, e.g., tablets

located in a cluster of a DFS can include indexed rowsets. Each rowset has rows and columns in which keys of arbitrary byte strings are stored. Each key has a value within an associated tablet 500. As illustrated in Fig. 5, rowsets 502, e.g., rowsets R1-R6, may be stored on a tablet. Each rowset is shown as an interval occupying one or more key spaces 504, based on the first and the last key. The keyspaces each span horizontally (e.g., x-axis) and may also be referred to by their primary sort key, represented as "Pk". For example, rowset R1 occupies keyspaces Pk1, while rowset R5 occupies keyspaces Pk1-Pk5. Thus, in this example, rowset R1 overlaps in range with rowset R5.

[0058] Utilizing the rowsets R1-R6 and associated keyspaces Pk1-Pk5, in some embodiments, a cost measure can be defined for performing a compaction on those rowsets. This cost measure can be applied to any rowset or rowsets, as a simple function based on known data, such as within a metadata file, which is maintained and updated on each tablet 500. The cost measure can be defined based on one or more variables within the tablet, e.g., as indicated via a hardware cost and/or a software cost associated with I/O operations. A first variable can be the widths of the selected rowsets in the tablet 500. A second variable can be the height of each keyspaces in the tablet 500.

[0059] The width of a rowset can be proportional to the percentage of a keyspaces that it spans. For example, the tablet 500 includes five keyspaces Pk1-Pk5. If each keyspaces is assumed to be equal to "w", then R1 has width 1w, R2 has width 1w, R3 has width 1w, R4 has width 2w, R5 has width 5w, and R6 has width 1w. In numerical terms, for example, each width "w" can represent, for example, 0.2 or 20% of the keyspaces spanned in the tablet 500. This width can be used to represent the probability that any read (or other operation) in a uniform random read workload will have to consult a particular rowset.

[0060] The height of a particular keyspaces can be defined as the number of rowsets that have overlapping key ranges. For example, Pk1 has a height of "2" because rowsets R1 and R5 overlap in Pk1. Keyspaces Pk2 has a height of "2" because rowsets R2 and R5 overlap in Pk2. Also, Pk3 has a height of "2", Pk 4 has a height of "3", and Pk5 has a height of "2". The height of a particular keyspaces can be used to represent the number of rowsets that will have to be consulted for a random read of that key. Thus, in Fig. 5, the height of the tablet ranges between two and three.

[0061] With these two variables, e.g., a width of a rowset in a keyspace and a number of rowsets included in a keyspace, a cost can be calculated specific to each operation performed on the tablet 500. For example, reads on a particular tablet (or rowset) may generally refer to accesses to data and may encompass specific operations such as inserts, random reads, short scans, long scans, and other operations known in the art. For inserts to be handled in the tablet 500, each rowset needs to be checked for a duplicate key of the known key value included in the insert request. Accordingly, the insert is made in the rowset having a keyspace including that key.

[0062] To more efficiently determine this result, by not querying each rowset, the rowset key ranges can be stored within an interval tree. Referring to Fig. 5, if, for example, data is to be inserted in a keyspace covered in the interval Pk4, then querying each rowset for a duplicate key would return the rowsets R4, R5 and R6. In some embodiments, consulting the interval tree can quickly return that R4, R5, and R6 may include the key corresponding to the insert without having to query rowsets R1 and R2. The associated cost measure to determine which of the three rowsets (R4, R5, or R6) includes the key for the insert is linear in those three rowsets. For example, the entire rowsets R4, R5 and R6 and keyspaces spanned by those rowsets will need to be queried to locate the corresponding key for the insert.

[0063] This cost measure can be represented in the following equation 1:

$$\text{Cost}_{\text{insert}} = n(C_{\text{bf}} + B_f * C_{\text{pk}}) \quad (\text{equation 1})$$

In equation 1, "n" can represent the height (i.e., the number of rowsets accessed during a scan) of the tablet at the keyspace including the given key, e.g., Pk4, the "C_{bf}" can represent the cost to run a bloom filter check on the rowsets R4, R5 and R6 in the keyspace, the "B_f" can represent the bloom filter false positive rate, and the "C_{pk}" can represent the cost of a primary key lookup (e.g., one seek in the interval tree).

[0064] In some embodiments, the bloom filter misses "B_f" equal to approximately one percent (1%) or lower, so the bloom filter checks dominates the cost measure equation and the second term (B_f*C_{pk}) can be ignored and would not drastically change the cost measure. In some embodiments, however, the primary key column (C_{pk}) is very large and

every primary key check incurs a disk seek. This causes the “ C_{pk} ” value to be orders of magnitude higher than “ C_{bf} ”, which is likely stored in the random access memory (RAM) element or solid-state drive (SSD). In the aforementioned embodiment, the term resulting from the bloom filter misses (B_f) would then be indispensable and would noticeably alter the cost measure. Accordingly, in embodiments in which a larger key space is to be accessed, the impact of the second term in the cost measure can be significant.

[0065] To process a random read, the cost measure approximates the cost measure for an insert. Accordingly, given a known key for the read, each potentially overlapping rowset in that key space is queried. Conversely, for short scans, bloom filters are not utilized. Short scans can include, for example, a seek with a primary key in one rowset retrieving one or more short data segments. Furthermore, in a short scan, the sequential I/O cost after finding the start key is small compared to the seek cost. For example, this comparison may include a ten millisecond (10ms) seek time and one megabyte (1MB) or less of sequential I/O.

[0066] For short scans, the cost measure is similar to the inserts and reads, except that all overlapping rowsets need to be looked up by the primary key. This can be represented in the following equation 2:

$$\text{Cost}_{\text{short}} = n * C_{pk} \quad (\text{equation 2})$$

Comparing equation 2 with equation 1, the only difference is that all bloom filter terms are removed, resulting in a direct relationship with the number of overlapping rowset in key column C_{pk} .

[0067] In the case of long scans, computing a cost measure for long scans requires additional terms to be factored into equation 2. For example, the disk space spanned by the scan and the disk bandwidth spanned by the scan are accounted for since, comparatively, a scan on 1MB (short) is going to be nominally faster than a scan on 10MB (long). In order to cover such a large amount of data, long scans can include full table scans or seeks on data from numerous rowsets. Factoring in the size of the rowsets in the long scan is therefore important to determine the cost measure, as represented in the following equation 3:

$$\text{Cost}_{\text{long}} = n \cdot C_{\text{pk}} + S/B_w \quad (\text{equation 3})$$

The additional term "S/B_w" accounts for the rowset size, where "S" can be the number of megabytes (MB) covered by the scan and "B_w" can be the disk bandwidth (MB/sec). The term "n" indicates the number of rowsets accessed during the scan.

[0068] The commonality shared among each of the aforementioned cost measure operations is that each rely significantly on the number of rowsets accessed in order to perform the operation, e.g., "n". Therefore, in accordance with embodiments disclosed herein, the cost of sequential I/Os to perform an operation on the log structured database directly correlates to the number of rowsets accessed. Based on this correlation, embodiments of the proposed compaction policy aim to minimize this cost by minimizing the number of rowsets accessed, while also taking into consideration that performing a read on a substantially large rowset causes increased latency for return.

[0069] Since minimizing the number of rowsets accessed during a read is not usually plausible, two strategies may be considered to account for this factor within the compaction policy. A first strategy includes merging rowsets with overlapping keyspaces in order to improve the cost for point queries, such as short scans, inserts and random reads, thereby reducing the overall height of the tablet. This strategy produces linear returns. The taller the tablet height becomes, the longer the return for one seek. In other words, the more number of rowsets on which to perform a seek, the longer it takes to return for one seek. A second strategy includes merging all rowsets to improve the ratio of sequential I/O to seeks, e.g., minimizing the number of rowsets, in order to improve the cost for long scans. In the latter strategy, the end result is an extremely large rowset (≥10MB) spanning an entire keyspace, assuming that the original rowsets are reasonably large. Accordingly, the cost is optimal, since any given access requires one seek and one bloom filter, at most. However, as mentioned previously, performing sequential I/O on a rowset this large results in diminished return time for one read. For example, one seek on 10MB of sequential I/O may take one hundred milliseconds (100ms) as opposed to ten milliseconds (10ms) for a seek on 1MB. Though either strategy may be implemented, the first strategy is therefore preferable over the second strategy since one factor, e.g., return

time, is not significantly diminished for the sake of the second factor, e.g., cost for performing a seek on the rowsets.

[0070] An additional consideration for implementing the second strategy is the cost of performing a compaction on such a large rowset and the wasted I/O on performing a compaction on all rowsets during each compaction in order to maintain the large rowset. Compacting such a large rowset also causes immense write amplification because each compaction rewrites the entire rowset. To summarize, in some embodiments, the efficiency gain of improving the ratio of sequential I/O to seeks by compacting the rowsets into a large rowset does not outweigh the diminished efficiency and increased cost of performing the compaction, nor the increased return time. Therefore, the first strategy of compacting overlapping rowsets can be the preferable approach.

[0071] Determining a candidate rowset on which to perform a compaction thus becomes a balancing act between the cost of performing a compaction on the rowset and the efficiency of the rowset prior to and after performing the compaction. In embodiments discussed herein, the cost of performing the compaction may be the sum of I/O performed by the compaction (e.g., Output (Input size)), assuming deletions are rare and the compaction inputs are large enough that sequential I/O outweighs any seeks required. By weighing these two factors of compaction cost and rowset efficiency against one another, the implemented compaction policy can selectively determine which rowsets should be compacted at any given time.

[0072] A final factor considered within the compaction policy is the way in which the compaction is performed. For example, performing a large compaction at longer time intervals decreases system performance during those compactions, whereas incrementally performing compactions more frequently maintains a consistent performance profile for the end-user applications. Incrementally performing the compactions on smaller data sets also allows the system to react more quickly to changes in the workload. For example, if an area of a keyspace becomes overloaded, the system is able to compact that area within a short time frame. This is because only one incremental compaction needs to occur prior to compacting the overloaded area rather than a very large and time consuming compaction.

[0073] To facilitate incremental compaction, the system may impart bounds on the data size for each compaction to read and write. For example, the bounds may include data sizes in the range of several hundred megabytes, which allows the compaction to occur in ten seconds (10s) or less. These bounds also help to maintain a consistent and lower cost for each compaction to occur in addition to a consistent time to complete the compaction.

[0074] Fig. 6 is an example diagram of numerous rowsets of bounded size stored in a plurality of keyspaces. In order for the system to limit a size of a rowset within predefined bounds, the DiskRowSet writer code may be modified to roll over to a new rowset once a threshold (e.g., upper bound) rowset size is met. In the example illustrated in Fig. 6, the maximum rowset size is 64MB, e.g., DiskRowSet max = 64MB.

[0075] When the new dataset stored in the MemRowSet 506 buffer is flushed to disk, this data is then distributed to a keyspace, e.g., Pk1, within the predefined bounds of DiskRowSet. For example, given a 200MB data flush from MemRowSet, R1 is 64MB, R2 is 64MB, R3 is 64MB, and R4 is the remaining 8MB. In some embodiments, the system may estimate the on-disk data size such that each rowset receives equal or approximately equal data sizes during the flush.

[0076] After several MemRowSet flushes, keyspaces often include various rowsets with small files spanning the entirety of the keyspace because these bounds cause roll over rowsets of lesser length. For example, in Fig. 6, R4 includes only 8MB of the 64MB maximum. These small sized rowset commonly occurs in, for example, uniform random insert load.

[0077] After several MemRowSet flushes, various smaller sized rowsets may be overlapping within a particular keyspace. For example, with respect to Fig. 6, Pk1 may additionally include rowset R5 (22MB) and rowset R6 (8MB). Rowset R5 and R6 can both overlap with rowset R4, causing the depth of the Pk1 to be three which, in turn, causes the average rowset depth in that keyspace to be three. Thus, each access to the tablet requires a query to each of the rowsets in that keyspace which decreases the performance of the tablet. Once the compaction policy is able to identify these rowsets to compact, then the output rowset, e.g., R7 (38MB) created from rowsets R4, R5, and R6 maybe

reorganized within the keyspace and include a single, non-overlapping rowset occupying the space of R4, assuming the range of R4 to be encompassing the ranges of rowsets R5 and R6. The average rowset depth in that keyspace then equals one. In the aforementioned embodiment, the improvement to the tablet performance outweighs the I/O cost of performing the compaction. The tablet configuration cost may then be optimized under a given I/O budget.

[0078] In order to facilitate identifying rowsets having overlapping key ranges for the compaction policy, the tablet code may be modified to store the rowsets in interval trees instead of simple lists. An interval may include a data structure which provides an efficient query for a set of intervals overlapping a given query point or query interval. Hence, any overlapping rowset within a particular keyspace may be identified within the tree.

[0079] Fig. 7 illustrates various overlapping rowsets within a keyspace Pk1 700, which may be identified by an interval tree (illustrated with the help of vertical dashed lines). In the embodiments illustrated in Figs. 7-9, each of the rowsets are assumed to be the same size, with overlapping portions appearing to span less range within the keyspace Pk1 700. However, with numerous overlapping rowsets in the keyspace, the efficiency of the system decreases because each access to the tablet queries each level of the rowsets, which equals the height of the keyspace. For example, in Fig. 7 the rowset depth, or keyspace height 702, ranges between two ($h=2$) and four ($h=4$).

[0080] As discussed previously, the height of a rowset being queried in a particular keyspace is linear to the cost of a single read or insert. Thus, to calculate the average cost of I/O operations on rowsets in a particular keyspace, the tablet height is integrated across that keyspace. This integration equates to summing the widths of all rowsets in that keyspace. For example, in Fig. 7, the cost ($R1+R2+R3+R4+R5$) equals approximately sixty-five (65). To decrease the height of the keyspace and, subsequently, the cost of performing an I/O operation, the compaction may be used to identify and select rowsets for compaction. These rowsets may be chosen based on the given I/O budget and the cost of optimal tablet configuration.

[0081] Fig. 8 illustrates an exemplary compaction of various rowsets in Pk1 800 shown in Fig. 7. For example, R1, R2, R4, may be selected for compaction in order to

reduce the height of the tablet by two. These rowsets may be selected as good candidates for compaction because each of the rowsets spans a smaller range (e.g., in comparison to the keyranges of compacting rowsets R1, R2, and R5) within key space Pk1 800 and is included in the maximum height range (e.g., 3) of the key space. A compaction process may be performed on the selected rowsets as described with reference to Fig. 4. Once compacted, the rowsets R1, R2, and R4 can map to rowsets R6, R7 and R8 in order to decrease depth in the tablet's rowsets, to remove overlapping rowsets, and create smaller sized rowsets. Each of these factors improve the performance and related I/O costs of the tablet. For example, the cost of performing I/O operations, such as inserts, is now reduced to forty (40). As mentioned previously, this is calculated by integrating the height of tablet across all rowsets. The compaction benefit for cost is therefore twenty-five (25) for a budget of three (3) I/O operations (e.g., 1 for seek and 2 for compaction process). To calculate the cost benefit, e.g., reduction, the cost of the compacted rowsets (e.g., sum of the output widths) is subtracted from the cost of the original rowsets (e.g., sum of the input widths). For example, the original cost (i.e., the sum of the input widths) in Fig. 7 is (65) and the output cost (i.e., the sum of the output widths) is (40), the difference of which equals (25). Additionally, because each of the output rowsets no longer overlap, the total width of those rowsets spans the union of the input rowset ranges. Thus, the aforementioned cost difference can also be computed as the cost of the original rowsets (e.g., sum of the input widths) minus the union width of the original rowsets.

[0082] It should also be noted that, though compacted, the total number of bytes have not changed in the tablet. Rather, the bytes have been rearranged in a more compact form and the height of the tablet is now equal to two (2). Additionally, it is noted that the aforementioned cost calculation may similarly be applied to string keys by treating the string data as unsigned integers, just as the key ranges within the embodiments described above are treated as integers to sum the widths of the rowsets.

[0083] Fig. 9 illustrates an alternate embodiment in which an incremental compaction is performed on the rowsets depicted in FIG. 7. Instead of selecting rowsets R1, R2, and R4 for compaction, the compaction policy may select rowsets R2, R4, and R5. By selecting R5, it can be assumed that at least one level of depth is removed from the tablet,

since the rowset spans the entire keyspace Pk1 900. Additionally, the selection of rowsets R2, R4, and R5 compacts each of the widest rowsets in the tablet, thereby reducing the range of keyspace spanned by each rowset in the tablet, which improves tablet performance for shorts scans, inserts, reads, etc.

[0084] After compacting rowsets R2, R4 and R5, the rowsets may be merged into rowsets R6, R7, and R8 as shown in Fig. 9. The height of the tablet is reduced to two (e.g., $h=2$), improving the performance of the tablet. Additionally, the cost of performing I/O operations is now reduced to thirty-five (35) based on the summed widths of each compacted rowset. Accordingly, the benefit of compaction is thirty (30), which is a slight improvement to the benefit in the embodiment illustrated in Fig. 8 while still operating under the same I/O budget of three (3). Accordingly, the compaction performed in Fig. 9 is the preferable option.

[0085] Fig. 10 includes a method for selecting rowsets on which to perform a compaction using the compaction policy described with reference to Figs. 7-9. The method may be implemented in an algorithm that includes instructions for a process executed on a server computer of the data node. The algorithm, for example, may be based on constant sized rowsets (e.g., bounded rowsets) as inputs and may be executed given a I/O budget for operations performed on the rowsets. The budget may include the number of I/O operations to perform the compaction on the selected rowsets. Accordingly, the algorithm for the method may include:

[0086] Given budget N rowsets:

For each pair of rowsets (A, B):

Evaluate BestForPair(A, B):

BestForPair(A, B):

Let union width = $\max(A.\text{max_key}, B.\text{max_key}) - \min(A.\text{min_key}, B.\text{min_key})$

Determine the subset R of rowsets that are fully included within the range A, B

Evaluate PickRowsetsWithBudget(R, N):

Set objective = $\text{sum}(\text{rowset width}) - \text{union width}$

If objective > best objective:

best solution = this set

[0087] An instance of the aforementioned algorithm may be implemented on each tablet in order to maintain an efficient tablet that provides optimal operations, such as reads, scans and inserts, at the lowest possible cost. The inputs include a number "N" of rowsets in a tablet being analyzed for compaction. Specifically, in step 1002, a set of two rowsets "A" and "B", i.e., a rowset pair, is analyzed in one iteration of the algorithm to determine if those rowsets should be selected for compaction. This rowset pair may be sequentially selected in the tablet and the compaction policy may be implemented continuously during each I/O in order to determine the rowset pair for the next compaction. Each step within the algorithm is further detailed in the following corresponding steps within the method illustrated in Fig. 10.

[0088] In step 1004, the union width is calculated for a pair of rowsets "A" and "B" being compared to determine the best rowset to compact within that pair. Accordingly, the union width, as discussed with reference to Fig. 9, is determined for the pair of rowsets by finding the maximum range of the key space spanned by the rowset pair. For example, in Fig. 7, the maximum key values of R1 and R2 are equal. However, the minimum key value for R1 is greater than the minimum key value of R2. Accordingly, the union is simply the range of key space P_{k1} covered by R2. This corresponds with the compaction performed on these rows in Fig. 8 in which the width of the compacted rows R6, R7 and R8 equals the width of R2.

[0089] In step 1006, a variable R, including a subset of the rowsets fully included within the range of key space covered by the rowset pair A and B, is determined. The subset R includes the combined rowsets in the key space within the key space range covered by A and B. Accordingly, referring back to Fig. 7, if A=R1 and B=R2, both R3 and R5 would not be included within the subset R because both include rowsets outside the range covered by R2. However, R4 includes rowsets fully included within the key space range covered by R2. Therefore, the rowsets within R1, R2 and R4 would be included in the subset R.

[0090] In step 1008, the subset of rowsets R is evaluated via PickRowsetsWithBudget over the input rowsets N to determine if the subset R provided the best rowset selection for compaction. This evaluation includes a sub-routine which calculates the cost change for variable R in N rowsets to determine which has the best cost change. For the highest pair, it is determined that the compaction policy selects that rowset pair for compaction because the cost benefit is worth the budgeted I/O. If the cost change of subset R is less than a cost change calculated for other pair of rowsets in N, then the budgeted I/O is not worth spending on the compaction of those rowsets. For example, calculating the "objective" for R with reference to Fig.7, where $A=R1$ and $B=R2$, would be $40-15=25$. Referencing the same Fig. 7, calculating the objective for R', where $A=R4$ and $B=R5$, would include $65-20=45$. Thus, in the aforementioned example, the rowsets in subset R' are the better rowset pair to compact. This is clearly demonstrated by the better cost change for the compaction performed in Fig. 9 than the cost change in Fig. 8.

[0091] PickRowsetWithBudget is evaluated in order to select the best rowsets to compact. For example, choosing the N rowsets in R which maximize the sum(rowset width) yields the highest objective. Thus, the PickRowsetWithBudget function can be solved by simply sorting the rowsets by their width and choosing the top N. For example, in Fig. 7, the rowsets having the largest widths were compacted in Fig. 9, illustrating that minimizing rowset range in the keyspaces provides an improved cost benefit.

[0092] The aforementioned method described with reference to Fig. 10 may also be extended to rowsets which are not a constant size. For example, while still maintaining bounds (e.g., a max threshold value) on the rowset size, many smaller rowsets may be stored in keyspaces due to more frequent flushes of smaller size datasets. For example, in Fig. 6, the smaller rowset R4 only includes 8MB of data and the flush included 64MB max size rowsets.

[0093] To properly account for these smaller rowset sizes while still complying with the aforementioned compaction policy, the given budget for performing the compaction is changed to a number of MB of I/O rather than the number of input rowsets (N). However, making this adjustment to the compaction policy also affects the PickRowsetWithBudget

function in which N is a variable. To handle this issue, a set of rowsets are chosen such that their entire file size falls within a budget and maximizes their total widths. This may be solved using a 0-1 knapsack problem solver in place of `PickRowsetWithBudget` in the algorithm referenced in Fig. 10. For example, the algorithm includes $O(n^2)$ calls to `BestForPair`. Each of these calls includes a 0-1 knapsack problem, which can result in a complexity of $O(n \cdot \text{max_budget})$. It logically follows then that the total complexity is cubic in the number of rowsets. Considering that any given tablet may have thousands of rowsets, the cubic complexity causes the pair selection process to become quite expensive. Thus, to optimize the approach to implementing the compaction policy on smaller rowsets, the order in which the each rowset pair (A, B) is considered is switched. For example, the previously stated algorithm can be modified as follows:

```
[0094] For each rowset A:
        candidates = all rowsets B such that B.min_key >= A.min_key
        sort candidates B by increasing B.max
        For each pair of rowsets (A, B):
            Evaluate BestForPair(A, B):
                BestForPair(A, B):
                    Let union width = max(A.max_key, B.max_key) - min(A.min_key,
                        B.min_key)
                    Determine the subset R of rowsets that are fully included within the
                        range A, B
                    Evaluate PickRowsetsWithBudget(R, N):
                    Set objective = sum(rowset width) - union width
                    If objective > best objective:
                        best solution = this set
```

[0095] In some embodiments of the compaction policy algorithm, changing the order in which the pairs are considered can simplify the `BestForPair` call. For example, because `B.max_key` does not decrease, the variable `R` in the step of "Determining the subset `R`" can remain identical to the previous call except that it includes `B` at the end. Thus, no additional looping is required. Additionally, the step of "Evaluating" including an

instantiation of the knapsack problem, is identical to the previous instantiation except the one additional item. Thus, this instantiation may be computed incrementally from the previous solution.

[0096] In some embodiments, upper bounds may also be calculated by solving the fractional knapsack problem and subsequently used to short circuit the more complex calculations. In some embodiments, the compaction policy is implemented on uniform workloads with equal sized rowsets. However, in additional embodiments, the compaction policy may be extended to non-uniform workloads, which may be experienced during use of a DFS. In practice, if a particular area of a key space is experiencing a higher workload, e.g., more accesses, those areas may be compacted prior to other areas in that keyspace. To implement compaction in those high workload areas, in some embodiments, the compaction policy algorithm may be modified to change all references to the "width" of a rowset to a cumulative distribution function (CDF). For example, $\text{width} = (\text{A.max_key}, \text{B.max_key}) - (\text{A.min_key}, \text{B.min_key})$ can be replaced by $\text{CDF}(\text{max_key}) - \text{CDF}(\text{min_key})$. The cumulative distribution function accounts for accesses over a lagging time window.

[0097] Fig. 11 illustrates a screenshot showing an example compaction policy log and a collection of exemplary rowsets. In a completely ideal scenario, upon compaction, data in a tablet would be arranged in a manner such that there would be a single rowset spanning a keyspace horizontally or multiple non-overlapping rowsets spanning a keyspace horizontally. But that would consume significant I/O operational cost associated with reading and re-writing all of the data in all the rowsets. Thus, the purpose of a compaction policy is to determine which rowsets to compact and when to compact so that more queries will hit fewer rowsets on an average. Thus, a compaction policy as disclosed herein, achieves a tradeoff of not doing compactions too often and yet doing enough compactions so as to achieve a reasonably compacted tablet. The disclosed compaction policy selects those rowsets in a collection of rowsets which when compacted would result in the improvement of the output metric (e.g., the average height of the tablet), given a fixed I/O budget. The speed (e.g., cost) of a query (e.g., an insert or an update) is linearly related to the number of rowsets that need to be crossed for a given key (e.g., the height of the table).

[0098] In Fig. 11, exemplary rowsets, their respective sizes, and their respective key ranges are shown as part of a compaction policy log in region 1152 of Fig. 11. These rowsets are also illustrated visually in region 1150 of Fig. 11. Rowset 85, rowset 71, rowset 41, and rowset 42 are indicated with the numerals 1102, 1104, 1106, and 1108 respectively. These rowsets are shown highlighted (dotted) in region 1150. As an intuitive example, it would result in a greater reduction in the height of the tablet in FIG. 11 if the rowsets in the left part of the tablet are compacted as opposed to the rowsets in the right side of the tablet. Accordingly, the disclosed compaction policy has selected the rowsets (illustrated with highlights) in the left to be compacted. The highlighted rowsets indicate those rowsets included in the collection of rowsets that would be compacted if a compaction were to run on the tablet. Thus, for a fixed (e.g., constrained) I/O budget, and considering the key ranges included in the rowsets, the size (e.g., in MB) of the rowsets, and other such parameters, the disclosed system has determined that the operational cost (e.g., resultant height of the tablet) of compacting the highlighted rowsets would result in the minimum operational cost (e.g., the minimum height) in comparison to compacting the other rowsets. Rowsets that are compacted are, on an average, those rowsets that overlap, are wide (e.g., span a wide keyrange) but are not too large in size. In some embodiments, the keyrange can be expressed as a difference of the maximum key value and the minimum key value in a rowset.

[0099] In some embodiments, the size of the rowsets are bounded by a predetermined maximum size (e.g., 32 MB or 128MB). Depending on the storage characteristics or the workload characteristics, the predetermined maximum can be changed. Sequential spinning disks spend a constant amount of time during a seek operation and a linear amount of time doing the I/O. If the predetermined maximum is very small, then a small amount of I/O would be necessary on a sequential spinning disk. Thus, a very small predetermined maximum size would be inefficient (e.g., low throughput for the drive) as it would result in wasted I/O because sequential spinning disks are faster for large amounts of I/O.

[00100] In some workloads however, a relatively small predetermined maximum size, e.g., a 32 MB limit can be chosen if the data is fine-grained in nature and the hardware is

more amenable to sequential I/O such as flash. An example of fine-grained data can be keyranges that include time-stamped data and such data is more likely to cause dynamically inserting recently-written data rather than inserting data that was written sometime back. Thus, there might be a sufficiently large amount of compaction activity (e.g., several incremental compactions) with respect to the recently-written data. Accordingly, if this maximum is chosen to be a very large number, then a large amount of data would be read and written every time compaction is performed. This might cause delays in insertion of the data (thereby slowing system performance) while compactions are being performed. This implies that several incremental compactions are redundant in such scenarios. Thus, a relatively small maximum size can be beneficial in adapting to rapid insertions of the data. For example, if the predetermined maximum size is 1 GB, most (or, even all) of the rowsets might be selected as candidates for compaction. However, if the workload is more static, and does not require too many insertions, then a single compaction of all the rowsets can be beneficial, and which can be achieved by a larger maximum size of the rowset.

[00101] Fig. 12 depicts an exemplary computer system architecture to perform one or more of the methodologies discussed herein. In the example of Fig. 12, the computer system 1200 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 1200 is intended to illustrate a hardware device on which any of the components depicted in the example of Fig. 1 (and any other components described in this specification) can be implemented. The computer system 1200 can be of any applicable known or convenient type. The components of the computer system 1200 can be coupled together via a bus or through some other known or convenient device.

[00102] 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.

[00103] 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.

[00104] 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 1200. 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.

[00105] 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.

[00106] 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.

[00107] In operation, the computer system 1200 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.

[00108] 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.

[00109] 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.

[00110] 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.

[00111] 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.

[00112] 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.

[00113] 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.

[00114] 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.

[00115] 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.

[00116] 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.

[00117] 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.

[00118] 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 sub-combinations. 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.

[00119] 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.

[00120] 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.

[00121] 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.

CLAIMS

What is claimed is:

1. A server-implemented method of selecting rowsets to compact in a tablet storing data associated with a distributed file system (DFS), the method comprising:
 - restricting, by a server, each rowset in a plurality of rowsets included in the tablet to have a size less than a predetermined size threshold, wherein the each rowset stores keys spanning a range that is less than or equal to the keyspace of the tablet;
 - determining a height of the tablet across the keyspace, wherein the height of the tablet is based on a number of rowsets having overlapping key ranges;
 - determining a rowset width of the each rowset in the keyspace of the tablet, wherein the rowset width of the each rowset is proportional to the percentage of the keyspace spanned by the rowset;
 - until a minimum operational cost is reached, iteratively calculating, by the server, an operational cost associated with compaction of two or more rowsets in the keyspace, wherein the cost is calculated by integrating the rowset width of the each rowset included in the tablet across the height of the tablet, wherein the operational cost depends on the size of the each rowset; and
 - selecting, by the server, the two or more rowsets for compaction based on the two or more rowsets resulting in the minimum operational cost over the each rowset included in the workspace.

2. The method of claim 1, further comprising:
 - receiving new data for the tablet in a memory element stored on a data node associated with the tablet; and
 - flushing the new data into a new rowset stored on the tablet.

3. The method of claim 2, further comprising:
if a new rowset reaches the size threshold, creating a new rowset in the keyspace to include any data overflow.
4. The method of claim 2, wherein the new data includes a first data and a second data, the new rowset is a first rowset, further comprising:
flushing the first data into the first rowset stored on the tablet;
flushing the second data into the second rowset stored on the tablet, wherein the first rowset and the second rowset are equally sized;
estimating an on-disk size for the first data and the second data prior to the flush, wherein the estimating is (i) based on the amount of the first data and the second data in the memory element and (ii) based on a range spanned by the first rowset and a range spanned by the second rowset.
5. The method of claim 4, further comprising:
sending a command to the data node to flush the first data and the second data in the memory element to the tablet, wherein the first rowset and the second rowset have non-overlapping key ranges.
6. The method of claim 1, further comprising:
translating the keyspace of the tablet into an interval tree for determining the height of the tablet based on an union of the overlapping key ranges.
7. The method of claim 6, further comprising:
determining a rowset height of the each rowset included in the tablet, wherein rowset heights and rowset widths of the two or more rowsets are greatest in the plurality of rowsets.
8. The method of claim 1, wherein the plurality of rowsets includes a first rowset and a second rowset with overlapping ranges, wherein the minimum operational cost

involves merging the first rowset and the second rowset resulting in a reduction in the height of the tablet, upon compaction and for a given input/output (I/O) budget.

9. The method of claim 8, wherein the operational cost is applicable for point queries, short scans, inserts, and random reads.

10. The method of claim 1, wherein the operational cost for long scans is directly proportional to a number of megabytes covered by a scan and inversely proportional to a disk bandwidth.

11. The method of claim 1, wherein the operational cost for long scans includes a direct proportionality to a cost of a key lookup and a number of rowsets accessed during the scan.

12. The method of claim 1, wherein the operational cost for short scans is directly proportional to a number of rowsets accessed during a scan and a cost of a key lookup, wherein overlapping rowsets are looked up by a respective primary key.

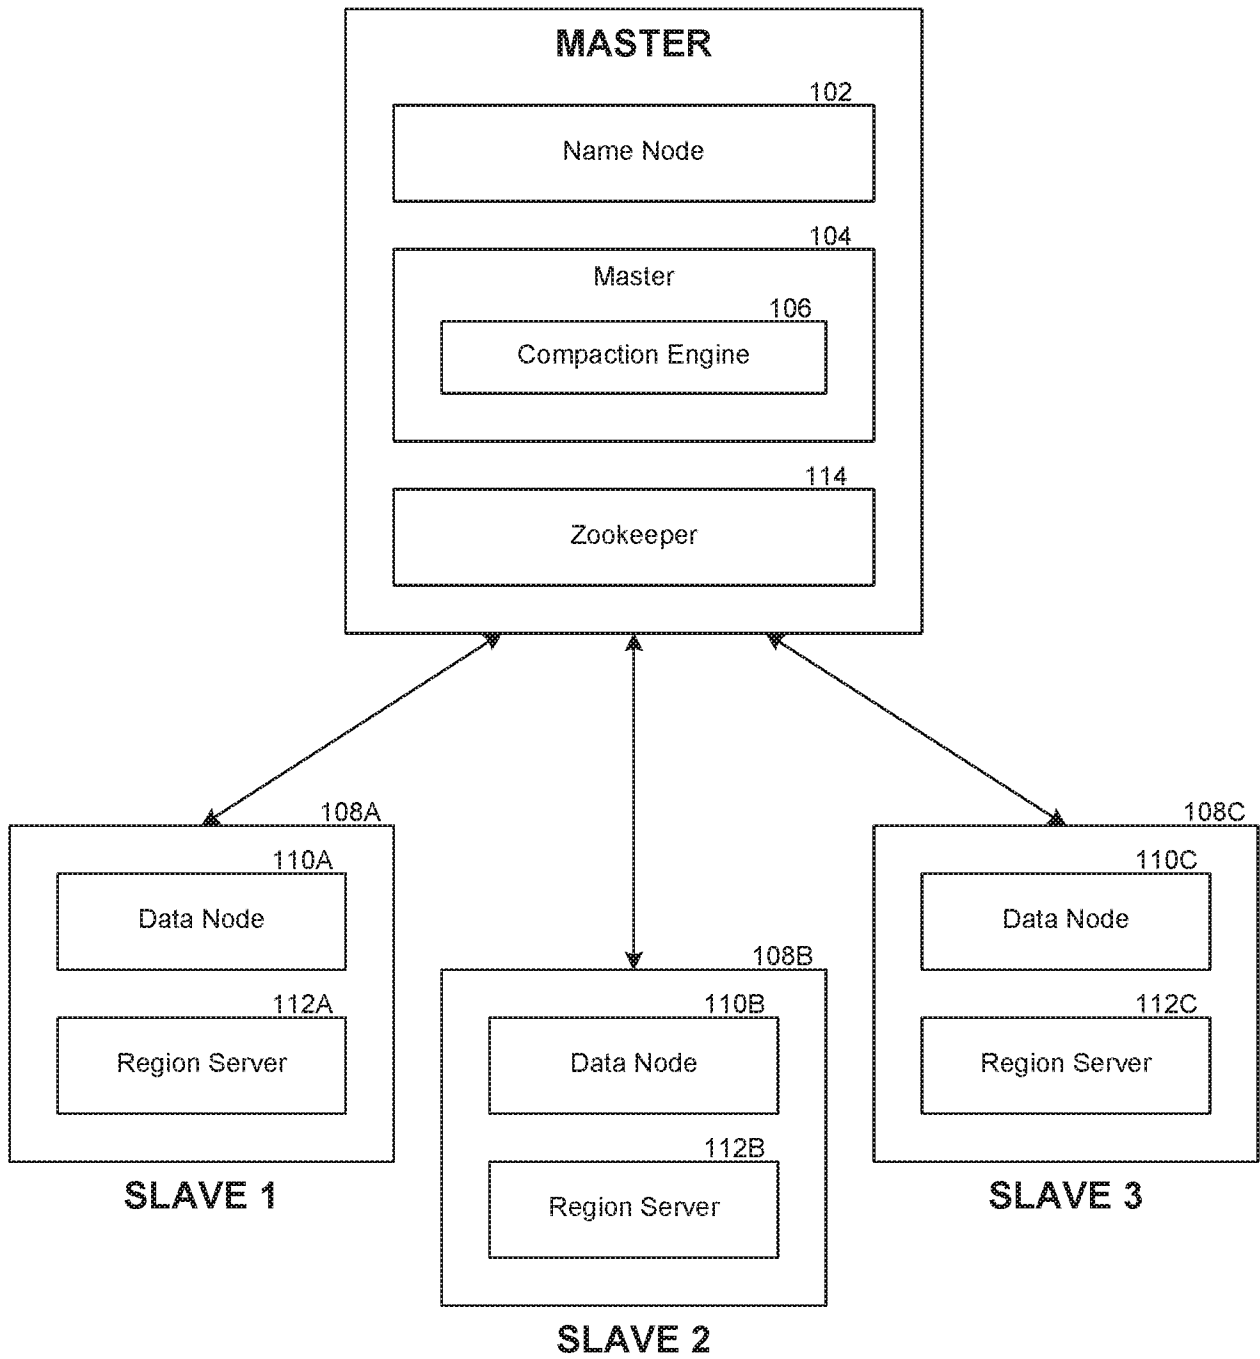


FIG. 1

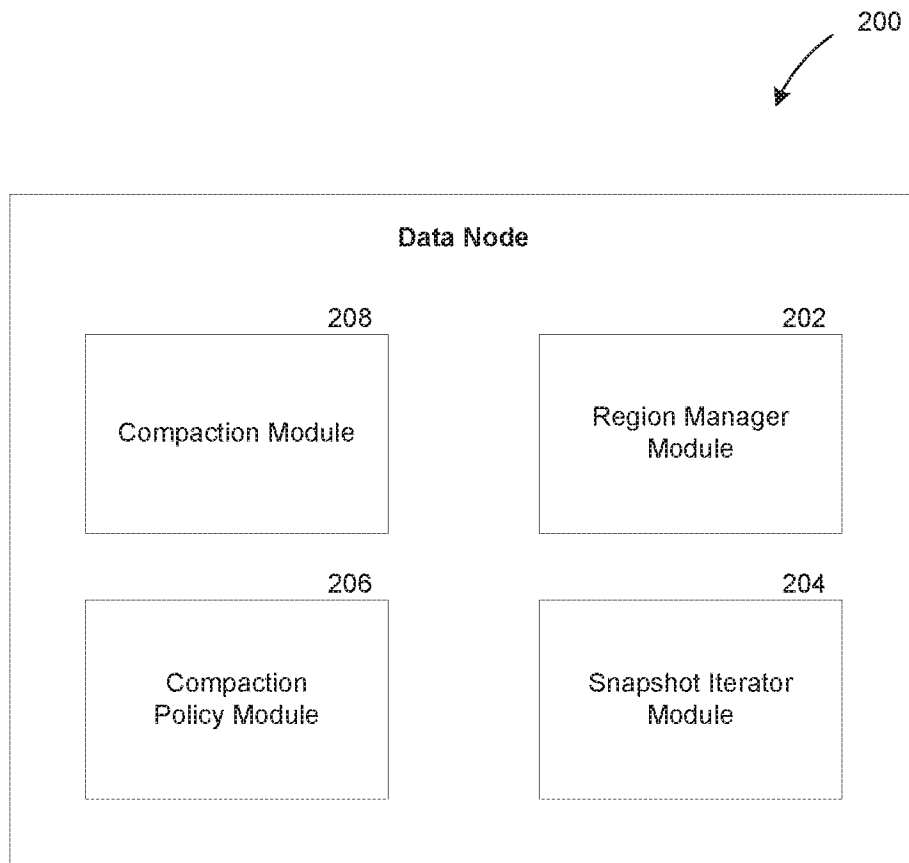


FIG. 2




FIG. 3A




FIG. 3B

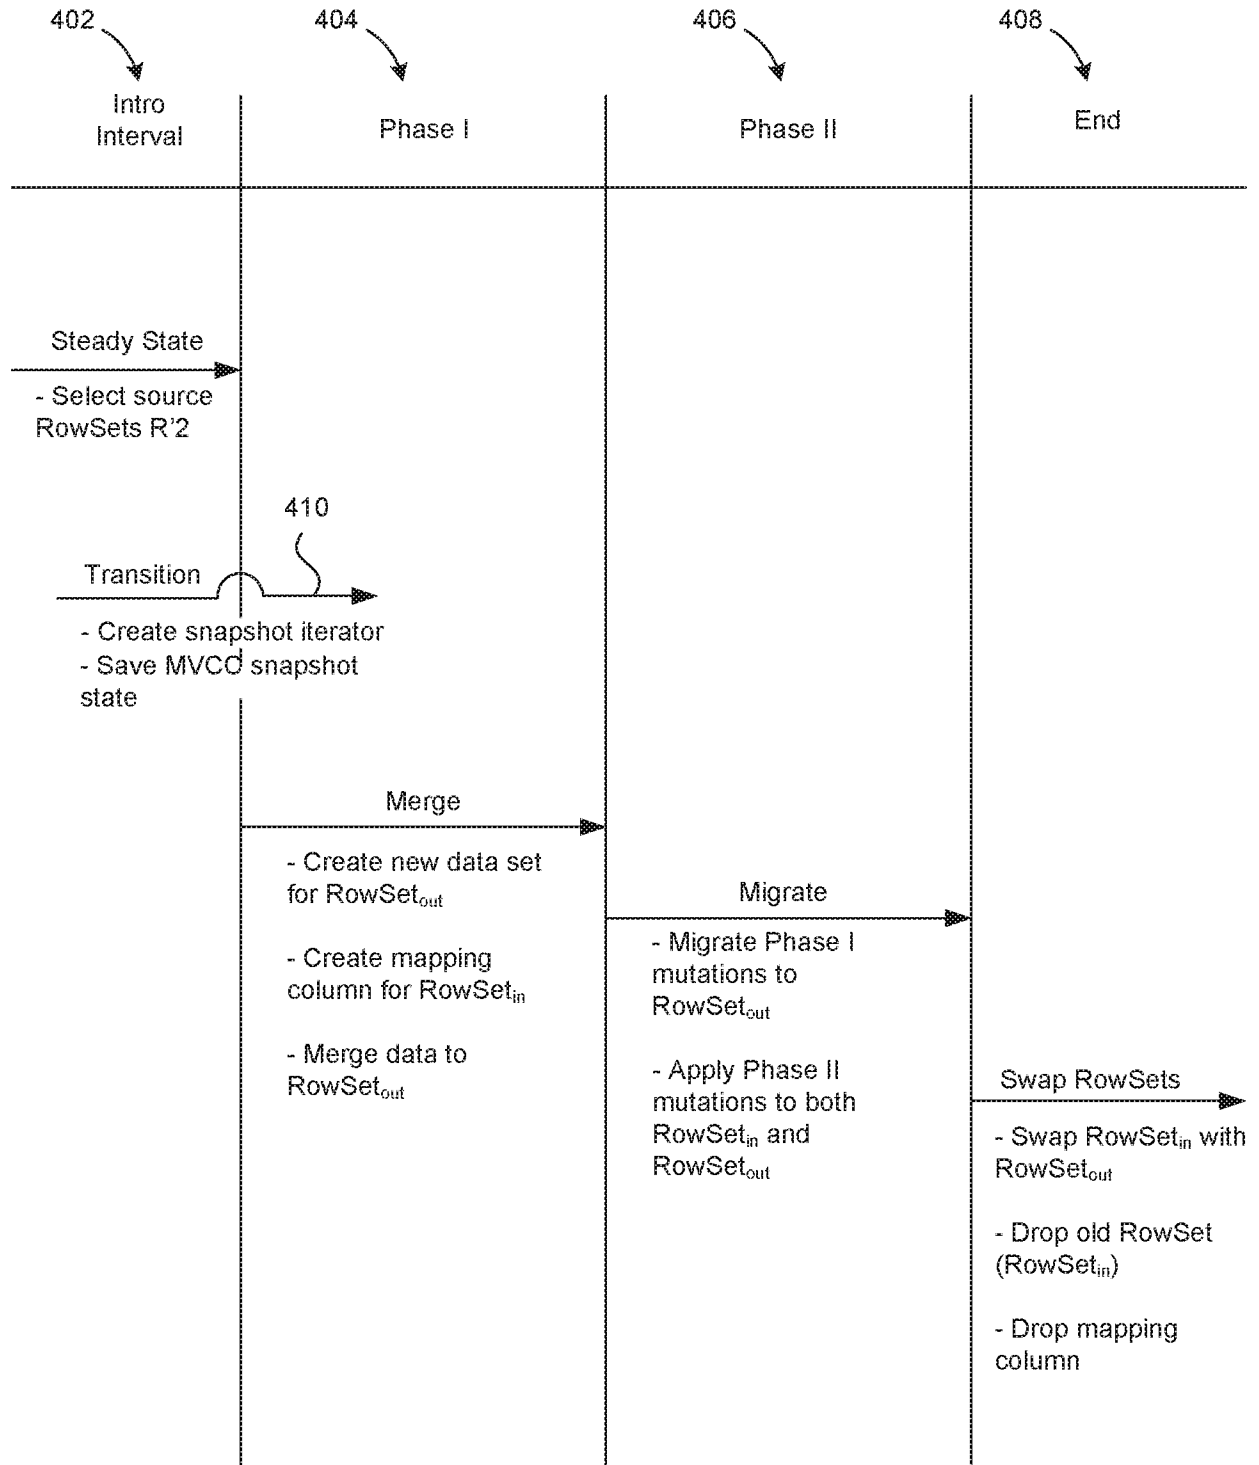


FIG. 4

5/9




FIG. 5




FIG. 6

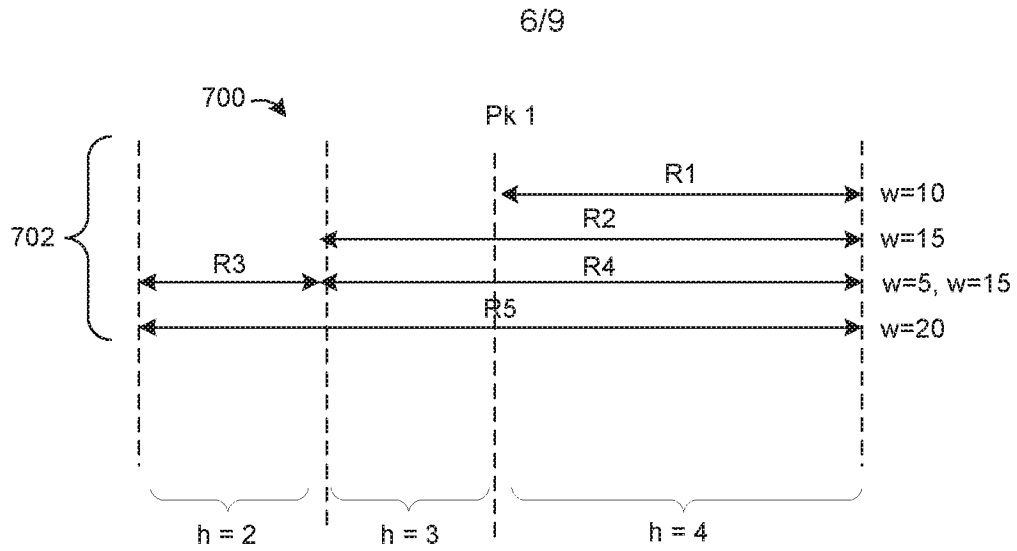


FIG. 7




FIG. 8


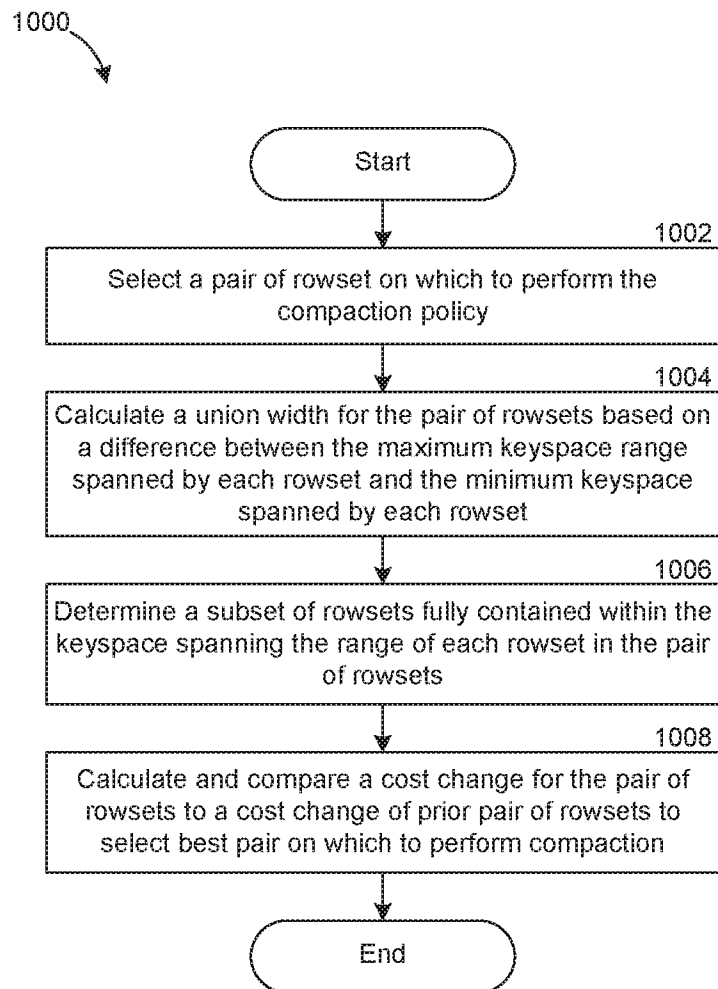


FIG. 9

7/9

**FIG. 10**

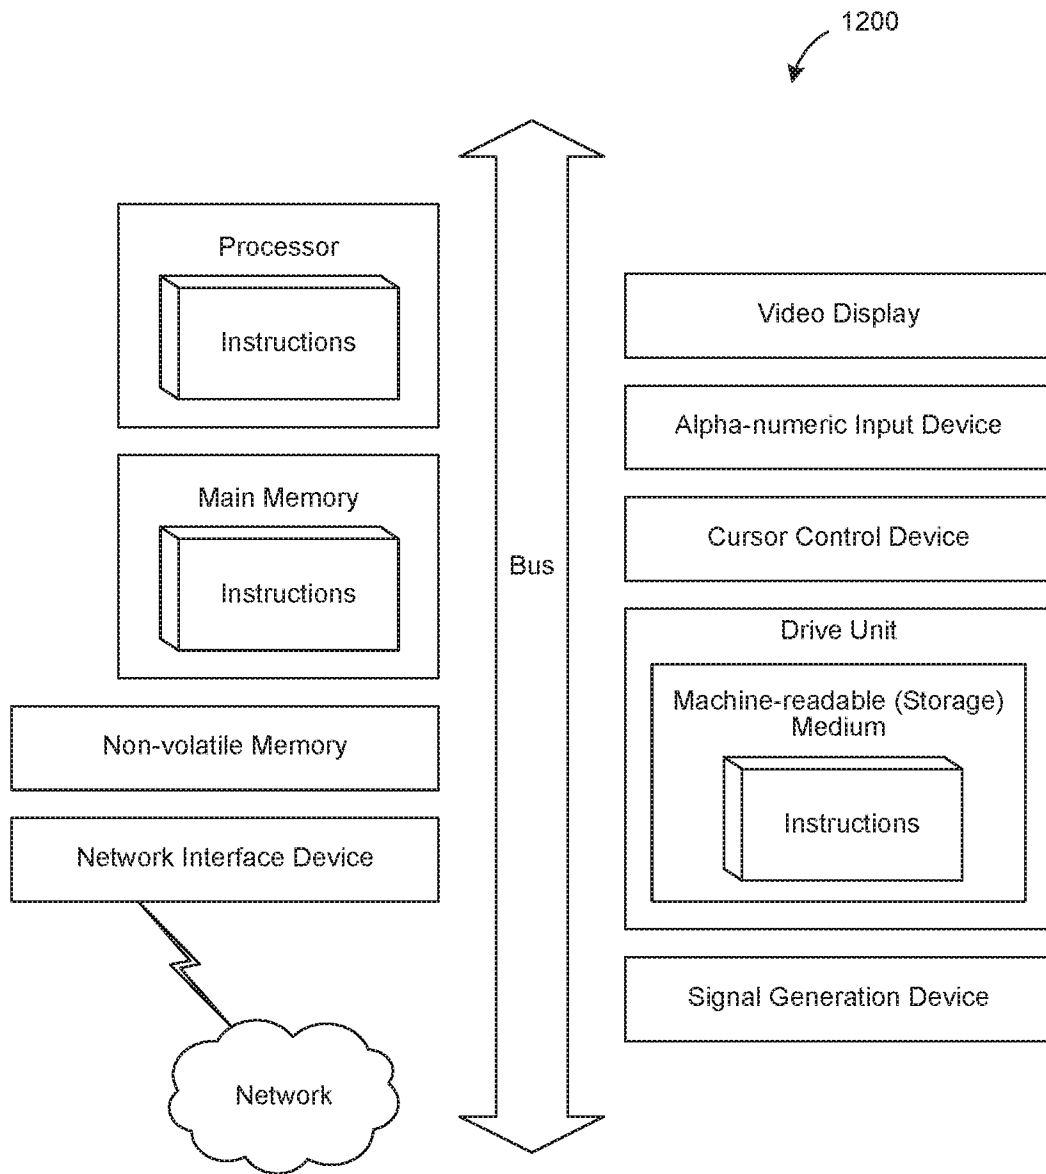


FIG. 12

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/022968**A. CLASSIFICATION OF SUBJECT MATTER****G06F 17/40(2006.01)i, G06F 3/06(2006.01)i, G06F 15/16(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 17/40; G06F 17/30; G06F 12/00; G06F 12/08; G06F 3/06; G06F 15/16Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: distributed file system, compaction policy, tablet, key space, rowset width, tablet height, minimum operational cost, and similar terms.**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2015-0026128 A1 (GOOGLE INC.) 22 January 2015 See paragraphs [0008], [0067], [0103], [0113], [0145], [0157], [0309], [0311], and [0320]; claim 6; and figures 12A-12B and 15A-15B.	1-12
A	US 2014-0297601 A1 (AVAYA INC.) 02 October 2014 See paragraphs [0007], [0009], [0033]-[0035], [0039], and [0042]; and figure 2.	1-12
A	US 2007-0033354 A1 (MICHAEL BURROWS et al.) 08 February 2007 See paragraphs [0004], [0017], [0026], [0041], and [0047]; and figure 1.	1-12
A	US 2012-0323971 A1 (KANTIKIRAN KRISHNA PASUPULETI) 20 December 2012 See paragraphs [0012], [0044], [0051], and [0061]-[0062].	1-12
A	US 2012-0072656 A1 (SHRIKAR ARCHAK et al.) 22 March 2012 See paragraphs [0007], [0025], [0035], [0039], and [0083].	1-12

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

07 July 2016 (07.07.2016)

Date of mailing of the international search report

08 July 2016 (08.07.2016)

Name and mailing address of the ISA/KR


International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/022968

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015-0026128 A1	22/01/2015	US 2011-0196833 A1 US 2011-0196900 A1 US 8868508 B2	11/08/2011 11/08/2011 21/10/2014
US 2014-0297601 A1	02/10/2014	GB 2512432 A US 9280570 B2	01/10/2014 08/03/2016
US 2007-0033354 A1	08/02/2007	EP 1917607 A2 EP 1983448 A2 EP 1983448 A3 EP 2434417 A1 US 7428524 B2 WO 2007-019007 A2 WO 2007-019007 A3	07/05/2008 22/10/2008 06/05/2009 28/03/2012 23/09/2008 15/02/2007 29/03/2007
US 2012-0323971 A1	20/12/2012	None	
US 2012-0072656 A1	22/03/2012	US 9355109 B2	31/05/2016