



- (51) **International Patent Classification:**
G06F 17/30 (2006.01) *G06F 12/14* (2006.01)
- (21) **International Application Number:**
PCT/US2013/041369
- (22) **International Filing Date:**
16 May 2013 (16.05.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant:** HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P. [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).
- (72) **Inventors:** CHANDRASEKHARAN, Kaushik; Longdown Avenue, Stoke Gifford, Bristol, Bristol BS34 8QZ (GB). CAMBLE, Peter Thomas; Longdown Avenue, Stoke Gifford, Bristol, Bristol BS34 8QZ (GB). TODD, Andrew; Longdown Avenue, Stoke Gifford, Bristol, Bris-

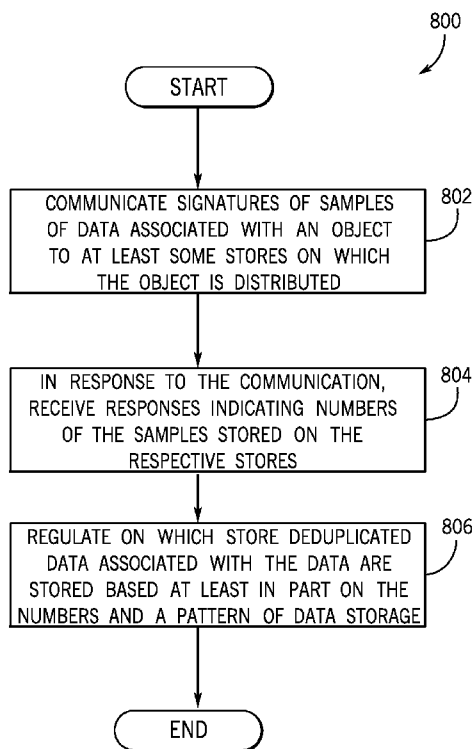
tol BS34 8QZ (GB). PELLY, Simon; Longdown Avenue, Stoke Gifford, Bristol, Bristol BS34 8QZ (GB). BHAGWAT, Deepavali Mahendra; 4209 Technology Drive, Fremont, California 94538 (US).

(74) **Agents:** ORTEGA, Arthur et al.; Hewlett-Packard Company, Intellectual Property Administration, Mail Stop 35 3404 E. Harmony Road, Fort Collins, Colorado 80528 (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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, 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, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ,

[Continued on next page]

(54) **Title:** SELECTING A STORE FOR DEDUPLICATED DATA



(57) **Abstract:** A technique includes communicating signatures of samples of first data associated with an object to at least some stores of a plurality of stores on which the object is distributed, and in response to the communication, receiving responses indicating numbers of the samples stored on the respective at least some of the plurality of stores and regulating on which store of the plurality of stores deduplicated data associated with the first data is stored based at least in part on the numbers and a pattern of data storage on the stores.

FIG. 8A

WO 2014/185916 A1

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, 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, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

Published:

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

SELECTING A STORE FOR DEDUPLICATED DATA

BACKGROUND

[0001] A typical computer network may have a backup and recovery system for purposes of restoring data (data contained in one or multiple files, for example) on the network to a prior state should the data become corrupted, be overwritten, subject to a viral attack, etc. The backup and recovery system typically includes mass storage devices, such as magnetic tape drives and/or hard drives; and the system may include physical and/or virtual removable storage devices.

[0002] For example, the backup and recovery system may store backup data on magnetic tapes, and after a transfer of backup data to a given magnetic tape, the tape may be removed from its tape drive and stored in a secure location, such as in a fireproof safe. The backup and recovery system may alternatively be a virtual tape library-based system that emulates and replaces the physical magnetic tape drive system. In this manner, with a virtual tape library-based system, virtual cartridges, instead of magnetic tapes, store the backup data.

BRIEF DESCRIPTION OF THE DRAWINGS

[0001] Fig. 1 is a schematic diagram of a computer network according to an example implementation.

[0002] Fig. 2 is an illustration of a data storage system according to an example implementation.

[0003] Figs. 3 and 4 are illustrations of a bidding process used by the data storage system of Fig. 2 to select a team member to receive deduplicated data according to an example implementation.

[0004] Fig. 5 is an illustration of the storage distribution of a teamed object across multiple team members according to an example implementation.

[0005] Fig. 6 is a flow diagram depicting a technique to store a teamed object on a cluster of team members according to an example implementation.

[0006] Figs. 7, 8A and 8B are flow diagrams depicting techniques to select team members for storage of deduplicated data according to example implementations.

[0007] Fig. 9 is a flow diagram depicting a technique to retrieve and report partial teamed object data according to an example implementation.

[0008] Fig. 10 is a table to illustrate the retrieval of teamed object data from team members according to an example implementation.

[0009] Fig. 11 is a flow diagram depicting a technique to distribute a master object manifest among team members according to an example implementation.

[0010] Fig. 12 is an illustration of the distribution of a master object manifest according to an example implementation.

[0011] Fig. 13 is an illustration of team member-controlled replication of duplicated data according to an example implementation.

[0012] Fig. 14 is an illustration of a non-hydrated replication process according to an example implementation.

DETAILED DESCRIPTION

[0013] Referring to Fig. 1, an example computer network 100 may include a backup and recovery system, which includes backup applications 132 and affiliated client applications 134 that execute on respective servers 110 (servers 110-1, 110-2. . .110Q, being depicted in Fig. 1 as examples). In this manner, from time to time, the backup application 132 identifies data to the affiliated client application 134 to be backed up on backup storage devices of the network 100. This data, in turn, is partitioned according to data containers called "objects" herein. From one backup session to the next, given objects that are stored on backup stores are created, deleted and modified. As disclosed herein, among its many functions discussed herein, the client application 134 is constructed to identify changes in the object data; select the stores on which the updated data are stored; and communicate the updated data to the selected stores.

[0014] The "stores" may be, as examples, independent computer systems or independent storage subsystems on the same computer system. For the specific example of Fig. 1, the stores are formed on respective nodes 150 (P nodes 150-1, 150-2. . .150P, being depicted in Fig. 1 as examples), which are coupled to the servers 110 via a network connection 140 (a local area network (LAN) connection, an Internet-based connection, a wide area network (WAN) connection, a combination of such connections, and so forth, depending on the particular implementation).

[0015] As disclosed herein, a given object is stored as a "teamed object" on a cluster, or group, of the stores. Due to the teamed nature, the "stores" are also referred to as "team members 170" herein. In this manner, in accordance with an example implementation, each team member 170 for a given "teamed object" stores "deduplicated data" for the object, where the deduplicated data are data formed from an initial set of data, along with data that represents the changes in the initially stored data. As such, deduplicated data may be retrieved from the team members 170 for a given teamed object to "rehydrate," or reconstruct, the object.

[0016] In general, the server 110 is a physical machine that is made of actual hardware 120 and actual machine executable instructions, or "software" 130. In this regard, a given server 110 may include such hardware 120, as one or multiple central processing units (CPUs) 122, a memory 124, a network interface 126, and so forth. In general, the memory 124 is a non-transitory memory, such as magnetic storage, optical storage, storage formed from semiconductor devices, and so forth. The memory 124 may store data locally for the server 110, as well as store instructions that when executed by the CPU(s) 122 cause the CPU(s) to provide one or more components of the machine executable instructions 130.

[0017] As illustrated in Fig. 1, the machine executable instructions 130 include the backup application 132 and the client application 134, as well as other possibly other applications that create, modify and delete objects.

[0018] A given team member 170 may be formed on a processing node 150 that is also an actual physical machine that is made of actual hardware 158 and actual machine executable instructions 159. The hardware 158 may include, as examples, one or multiple central processing units (CPUs) 160, a network interface and a memory 162. The memory 162 is a non-transitory memory and may be a magnetic storage-based memory, an optical storage-based memory, a semiconductor storage-based memory, and so forth, depending on the particular implementation. The node 150 may include machine executable instructions 159 that include, for example, a team member client application 168.

[0019] Thus, a cluster, or group, of team members 170 existing on potentially multiple nodes 150 may form the storage for a given teamed object, in accordance with an example implementation. Moreover, although described herein as existing on separate nodes 150, in accordance with further implementations, a given teamed object may be stored on independent team members, in which two or more the team members are located on the same

node 150. Thus, many implementations are contemplated, which are within the scope of the appended claims.

[0020] Referring to Fig. 2 in conjunction with Fig. 1, in accordance with an example implementation, a teamed object 208 may be presented as a single logical object to a given client application 134, although data for the object 208 are distributed over a group, or cluster, of team members 170. This logical presentation of the teamed object provides applications a single pool of storage, which spans the otherwise independent pools of storage within the cluster.

[0021] The client application 134, in accordance with some implementations, does not store locally any information regarding the contents of a given teamed object. This allows multiple client applications 134 and therefore, multiple backup applications 132, to access the same teamed object simultaneously and also avoid creating dependencies between specific client applications 134 and the data stored.

[0022] As disclosed further herein, to simplify integration with existing backup applications, each team member 170 may be aware of the other team members 170 for a given teamed object and may instruct the client application 134 of their locations. This allows the backup application 132 to connect to any one of the team members 170 and further allows the client application 134 to silently open up connections with all of the team members 170. This may help avoid exposing complex configurations and allow management of teamed objects within the backup applications 132, which may be designed, for example, with single end point topologies.

[0023] Because a given teamed object involves multiple independently operating team members 170, in accordance with some implementations, the client application 134 consolidates information from the multiple team members 170 into meaningful information that is communicated to the affiliated backup application 132. In this manner, in accordance with some implementations, a given team member 170 may store individual lists, data

job records, copy job records, and so forth, which a given client application 134 may consolidate into meaningful information for the backup application 132. For many of the fields that are returned in the list, the client application 134 provides a summation of all the returned values from across the team member stores, for example, the amount of user data stored. For other fields, the client application 134 may "wrap up" the individual statuses into an overall status, such as the most severe state(s), for example.

[0024] When the client application 134 performs a list operation across each of the team members 170, the absolute order of the entries in the list is not guaranteed. For example, two teamed objects may be created virtually simultaneously, and for a given team member 170, object one may be stored first, whereas on another given other team member 170 object two may have been created first. For purposes of providing a single, stable list to the backup application 132, universal identifications are used and list operations are used to search across the storage within a reasonable time window looking for associated records. By having a time-based window over which the search is run, a situation may be avoided in which the entire database is searched on each time member 170 looking for entries, which may under certain circumstances not even exist. As an example, the time window may be a time window in the range of approximately five minutes, but this window may vary, depending on the particular configuration and/or implementation.

[0025] As depicted in Fig. 2, although a given distributed teamed object 208 may be distributed among multiple team members 170, a given client application 134 has the logical view of a single corresponding object, which allows each team member 170 to potentially be executing a different release of the application software. This allows users to perform rolling updates to their software without having to be concerned about maintaining consistent software versions across all of the team members 170. In cases in which the capabilities of the software differ among the team members 170, the client application 134 assumes the least capability from across the team members 170.

[0026] In accordance with example implementations, for purposes of achieving acceptable deduplication performance across multiple independent team members 170, a bidding technique is used, with a goal of sending similar data to the same team member 170 and load balancing new, unmatched data across the remaining team members 170. In this bidding technique, for a given unit of data (a data segment, for example) the client application 134 requests bids from the team members 170, receives corresponding bid responses from the team members 170, selects one of the team members 170 based on the responses and communicates deduplicated data for the unit of data to the selected team member 170. Moreover, as further described above, the client application 134 may regulate when bidding is and is not used.

[0027] As a more specific example, referring to Fig. 3 in conjunction with Fig. 1, in accordance with some implementations, the client application 134 processes incoming data 310 to be backed up for purposes of loading balancing the storage of the data according to the illustration 300 of Fig. 3. The client application 134 receives the incoming data 310 in a buffer 312. In this manner, the buffer 312 stores a given data segment 314 or multiple data segments 314, depending on the particular implementation. Regardless of the storage, each data segment 314 is processed in the following manner.

[0028] A chunking (or chunk) module 316 of the client application 134 transforms the data segment 314 into corresponding chunks 318. For example, in accordance with some implementations, the chunking module 316 may apply a two thresholds to divisors (TTTTD) variable chunking algorithm, which produces an average chunk of approximately four kilobytes (kB). Other chunking algorithms may be used, in accordance with other implementations. In general, the chunking algorithm may enhance the likelihood of isolating identical chunks within successive backups, where the absolute location of the chunk may have moved.

[0029] Next, the client application 134 uses a hashing (or hash) module 320 to determine corresponding digital signatures, or hashes 324, for the chunks 318. As an example, in accordance with example implementations, the hashing module 320 may generate an SHA-1 hash for each chunk 318, although other hashes may be used, in accordance with further implementations. In general, a given hash 324 serves as a reasonably certain "fingerprint" for the associated chunk 318; and, on average, the size of the hash 324 may be relatively small, as compared to the size of the chunk 318 (approximately 0.5% of the size of the chunk 318, for example). Therefore, the hash 324 permits a relatively easy, low bandwidth way to identify an associated chunk 318. In accordance with example implementations, the hash module 320 determines hashes 324 for the corresponding chunks 318 and sends or makes available this list of hashes 324 to a bidding (or bid) module 330 of the client application 134.

[0030] The bidding module 330, in accordance with example implementations, communicates 334 a sparse index of the hashes to each of the team members 170 for the teamed object. In other words, in accordance with some implementations, the bidding module 330 communicates a statistically representative set of samples of the hashes 324 to the team members 170. Fig. 3 illustrates one such example for team member 170 that receives a sparse index of hashes.

[0031] It is noted that in accordance with some implementations, the bidding module 330 may communicate all of the hashes 324 for a given data segment 314 to each team member 170. However, in accordance with an example implementation, a single, relatively large list of hashes for matching may not be practical. In this regard, a twenty byte SHA-1 hash for each average size chunk of 4 kB means 5 gigabytes (GB) of memory for each one terabyte (TB) of unique data. To the contrary, the sparse hash index 334 contains a statistically chosen subset of hashes, such that these hashes adequately represent the chunks 318 while collectively being of a significantly smaller size (between 1 to 10 percent of the size of all of the hashes, for example).

When a portion of hashes is communicated to a given team member 170 for matching, the team member 170 determines the number of corresponding hashes that match.

[0032] In this manner, in accordance with some implementations, each team member 170 assesses matches and responds to the sparse index as follows. The team member 170 includes a bid matching (or match) module 370 that compares the sparse index against a list of hashes of the data stored in the team member 170. A successful sparse index match may be referred to as a "hook" because the sparse index is held in random access memory (RAM), for example, of the team member 170, and as such, the sparse index lookup may be relatively "cheap," in terms of system resources.

[0033] If the bid matching module 370 identifies one or more hooks in the sparse index, the module 370 may then, in accordance with example implementations, perform a more detailed matching, involving reading on-disk manifests pointed to by the sparse index hooks. Because this latter step involves disk seeks, which are relatively slow, this may be a relatively expensive process. To mitigate the use of the disk seeks, in accordance with example implementations, the on-disk manifests are read in some priority order based on the expected number of extra matches that will be found, with some stopping condition applied when there are many hooks, to keep performance up at the expense of a relatively small reduction in deduplication.

[0034] By providing the client application 134 a method of querying the sparse hash index, the sparse index hook count may be used to determine the probability of a given team member 170 matching the chunks 318.

[0035] In addition to the sparse index hook count, other information about the team member 170, such as the storage capacity and storage utilization (as two examples) may be communicated back to the client application 134 as part of the bid response. This information may then be used by the client application 134 to make a decision about which team member 170 to select

for purposes of routing all of the remaining hashes and the subsequent deduplicated data for the segment 314.

[0036] Thus, in accordance with some implementations, the bidding involves the bidding module 330 of the client application 134 communicating 334 the sparse index of hashes to the bidding match module 370 of each team member 170. The bidding match module 370 then communicates a bid response 374 to a router (or route) module 340 of the client application 134. As an example, in accordance with example implementations, the router module 340 may receive 374 one or more matches from the bidding match module 370. The router module 340 determines, based on the similar responses from the team members 170, which team member 170 is to receive the deduplicated chunks 318 of data for the segment 314.

[0037] After the router module 340 has selected the particular team member 170 (assumed for this example to be the team member 170 of Fig. 3), the router module 340 communicates, or sends 342, all of the remaining hashes for the chunks 318 of the data segment 314 to a matching (or match) module 380 of the team member 170. The matching module 380 compares all of the hashes of the chunks 318 to the corresponding hashes of data stored on the team member 170. The matching module 380 communicates 384 the matches to a compression (or compress) module 344 of the client application 134. In this manner, the matches inform the compression module 344 as to the unique chunks 318, i.e., the chunks 318 that are not stored on the team member 170. In response, the compression module 344 performs deduplication to communicate, or send 350, the unique chunks (e.g., the chunks of changed data) to a storage module 390 of the team member 170, which commits the new chunks to the team member 170.

[0038] The purpose of regular bidding is to route similar data chunks to the same team member 170. It is noted that each time a decision is made to change the team member 170, the segments that are routed to the new team members may have a negative impact on the overall deduplication ratio. This

is due to the relatively high likelihood that the data segment boundary does not align with the deduplication segment boundaries, and therefore, some duplicated data may be stored again. The fragmentation of the data stream may therefore be something that is controlled in a manner to minimize the reduction of the deduplication ratio, in a process that is further described below.

[0039] Fig. 4 depicts an illustration 400 of the communications that occur when a given team member 170 is not selected in the bidding process. In this regard, in response to the communication 334 of the sparse hash index, the bid matching module 370 of the team member 170 communicates a bid response 374, which for this example is not a winning bid. Therefore, the router module 340, for this example, sends, or communicates 410, a skip message to a skip module 420 of the team member 170, thereby informing the team member 170 of the bypassing of the member 170 for this particular data segment 314.

[0040] Referring to Fig. 5, in conjunction with Fig. 1, when the backup application 132 creates a given teamed object 500, the client application 134 causes corresponding objects 520 (objects 520-1, 520-2, 520-3, 520-4, being depicted in Fig. 5 as examples) to be stored on corresponding team members 170. Data items 318 of the teamed object 500 are distributed among the team members 170 to form the corresponding objects 520. For example, in the teamed object 500, a data item 318-A of the teamed object 500 corresponds to data item 318-A of object 520-1, whereas data item 318-F of the teamed object 500 corresponds data item 318-F of the object 520-4. As can be seen from Fig. 5, the data for a given teamed object may be distributed on a given team member 170 in an order different from the order in which the data appears in the teamed object, as the local ordering is left up to the individual team members 170, in accordance with an example implementation. As depicted in Fig. 5, each object 520 may contain data 550 that is not part of the teamed store. Therefore, the team members 170 may

track regions, which contain data for a given teamed store and regions for data that are not part of the given teamed store.

[0041] Thus, referring to Fig. 6, in accordance with an example implementation, a technique 600 includes communicating (block 602) chunk hashes (a list of chunk hashes, for example) to multiple team members, or stores, and receiving (block 604) responses from the storage, indicating a distribution of associated chunks in the stores. As described above, in example implementations, the list may be a sparse list of hashes. A store is selected (block 606) based at least in part on the responses, and deduplicated data are communicated to the selected store, pursuant to block 608.

[0042] In accordance with some implementations, in order for a team member 170 to be considered for a winning bid, the team member 170 first satisfies the criteria of matching a certain number key hashes above a certain threshold. In this manner, such a technique defaults to routing data to a "sticky" team member, i.e., the routing "sticks" to a "sticky team member" until the threshold is surpassed, in accordance with example implementations. By remaining, or sticking, with a team member 170 for several data segments 314 when matches are not above a certain threshold, many time contiguous segments with predominately new data (called "seeding data" herein) are routed to the same team member 170. Seeding large contiguous regions to the same team member 170 may help improve the overall deduplication ratio in a future backup. This is because for a future backup, the backup stream contents may vary to a degree and hence the segments may be aligned differently.

[0043] Therefore, if the segments are seeded to a different team member 170 for each data segment 314, segments 314 in a subsequent backup stream may have chunks straddling two team members 170. With a relatively long contiguous sticky region, deduplication may be lost at the ends of the region, not at each segment boundary within the region (as all data in that region was

stored on the same team member 170). Thus, a high deduplication ratio may be expected if all the data were routed in a backup to a single team member 170. However, such a technique may not aid in capacity balancing across the team members 170. Therefore, the sticky threshold may be selected to be small enough to be able to "stick to" another team member 170 often enough to seed across all team member 170, but the sticky threshold is large enough to keep the future deduplication ratio relatively high.

[0044] Thus, referring to Fig. 7, in general, a technique 700 may be employed, in accordance with example implementations. Pursuant to the technique 700, a list of chunk hashes is communicated (block 702) to multiple team members 170, or stores. Responses are received (block 704) from the stores, where each response indicates a number of matches. The technique includes selecting (block 706) a store based at least in part on a comparison of a match of the numbers to a threshold, such as the "sticky threshold" mentioned above. The deduplicated data are then communicated (block 708) to the selected store.

[0045] The deduplication may be performed between backups from the same system rather than between systems, so that when a system is first backed up, a considerable amount of chunk data may be stored from that system. The first backup is referred to as "seeding" herein and the initial data are referred to as "seeding data" herein. For purposes of avoiding excessive region fragmentation during seeding, a "sticky routing" technique may be used. In general, sticky routing attempts to stripe seeding data across the team members 170 in relatively large contiguous regions (regions on the order of tens of gigabytes (GB), for example), but the technique still routes data segments to other team members 170, if the team members 170 may deduplicate them well enough.

[0046] For a given data segment bid, if no team member 170 has a hook match count above a given threshold (called the "bid threshold" herein), then there is no "bid winner." This threshold may be a fixed threshold or may be a

threshold that is varied based on feedback obtained during the backup. If there is no bid winner, then the corresponding data segment contains seeding data, so that the data segment is routed to the current seeding team member 170.

[0047] In accordance with some implementations, at the beginning of the backup, the client application 134 may select the initial seeding team member 170 using a random or pseudo random technique. This may avoid a situation in which a teamed store is created and all of the first night's backups, starting at the same time, are seeded to the same team member 170. After a fixed amount of data is written, a new seeding team member 170 may be selected based on capacity utilization (the team member 170 having the most free disk space, for example). This technique levels disk usage across the team members 170, as the application 134 stores the seeding data.

[0048] Referring to Fig. 8A, a technique 800 in accordance with an example implementation includes communicating (block 802) signatures of samples of data associated with an object to at least some stores; and in response to the communication, receiving (block 804) responses indicating numbers of the samples stored on the respective stores. The technique 800 further includes regulating (block 806) on which store deduplicated data associated with the first data are stored based at least in part on the numbers and a pattern of data storage on the stores.

[0049] As a more specific example, Fig. 8B depicts a technique 850 in accordance with an example implementation. Pursuant to the technique 850, a list of chunk hashes is communicated (block 852) to multiple team members, or stores; and responses are received (block 854) from the stores, where each response indicates a number of matches. The technique 850 includes determining (decision block 860) whether the bid threshold has been exceeded. If not, then the data segment is seeding data, and the data segment is communicated (block 862) to the current seeding team member 170.

[0050] Otherwise, if a determination is made (decision block 860) that the bid threshold has been exceeded, the technique 850 includes determining (decision block 864) whether the current bid winner is the same bid winner as the immediate previous bid winner. If so and if the bid winner is a team member other than the currently selected team member (as determined in decision block 868), then a re-routing occurs and the data segment is routed to the current bid winner, pursuant to block 870. Otherwise, if in decision block 864 a determination is made that the current bid is not the same as the previous bid winner or if a determination is made, pursuant to decision block 868, that re-routing is not to occur, then the data is communicated to the currently selected team member without re-routing, pursuant to block 866.

[0051] In accordance with further implementations, the client application 134 may selectively suspend the bidding (and the communication of the hashes) based at least in part on a prediction of future bidding activity. For example, the client application 134 may predict when a region of "flux" exists in which time contiguous data segments 314 would end up being routed to different team members 170 if bids were made for these data segments 314. The client application 134 may temporarily suspend the bidding process when the application 134 predicts a region of flux, in accordance with example implementations.

[0052] For example, in accordance with example implementations, a region of flux may be predicted based on the number of measurable factors and/or statistics in a historic window for the current backup session. The factors may include measurements of such criteria as the number of times the bid winner has previously changed, the number of matches seen with the bid losers, the number of matches seen with the bid winner, and the amount of data written to the current bid winner. Using a calculated probability derived from these measurements, the client application 134 may elect not to perform a bid operation for a certain number of time consecutive data segments 314 and instead continue to route data segments 314 to the current winner without performing bidding for these segments 314.

[0053] In accordance with example implementations, all access to a teamed object is performed in a command or data session basis using a client-side code library of the application 134. The client library may be given the address of any team member 170 in the teamed store, connect to it and find the addresses of all the other team members 170. The client library may connect to the other team members 170 as well, thereby establishing the command or data session. All team members may not, however, be available for a given session. The team members 170 to which connections were successfully connected in this session may be reported back to the user of the client application library, so that the user may decide whether the user wants to continue with the session.

[0054] The client application 134 serves as an aggregator of information that is stored in/retrieved from the team members 170. By allowing sessions to be established with a subset of team members 170, the user of the client application 134 library is presented with a view (via a graphical user interface (GUI) 136 (see Fig. 1), for example) detailing a subset of the information that is available across the team members 170.

[0055] For example, when listing a given teamed object, the teamed object may have been created (and partitioned) across team members A, B and C, as those team members may have been, for example, the team members that were available at the time the command session was opened. If a list of that teamed object on a command session open to team members B, C and D is created, then the information available for the object in team members B and C are aggregated and presented to the client library, with the information for team member A not being presented.

[0056] When listing the teamed object, the client application 134 reports which team members 170 the team member was created on and last modified on. If the set of team members 170 for which the current command session is opened is not the same as the set on which the object was created and the set on which it was last modified, the client application 134 highlights to the

user that an incomplete view of the object is being presented. With this information, the user may decide how to interpret the listing.

[0057] Thus, referring to Fig. 9, in accordance with an example implementation, a technique 900 includes attempting (block 902) to open communication with all team members 170 that collectively store data for a given distributed teamed object in response to a request to access the object. If a decision is made (decision block 904) that all team members 170 are not present in the session for which the object was created and modified, the technique 900 includes noting (block 906) the absent team member(s). The available chunks for the teamed object are then retrieved, pursuant to block 908. If a determination is made (decision block 910) that degraded information is being reported, then the technique 900 includes reporting (block 912) information about the degraded state of the retrieved data, including identifying how the object is incomplete. With this degraded information, if any, the results are reported, pursuant to block 914.

[0058] In accordance with example implementations, the client application 134 also uses this information when aggregating the listing of multiple teamed objects. The challenge relates to how to present a page of listing results to the user by stitching together pages of results from the team members 170, effectively being a windowed multi-way merge. To perform this in an efficient manner, the client application 134 minimizes the number of pages of results retrieved from each team member for each page of results presented to the user.

[0059] In particular, the client application 134 uses the following three items of information it receives from each team member 170, in accordance with example implementations: 1.) a team-wide unique identification (ID) for each teamed object (or teamed data job), so that records returned from each team member 170 relate to the same teamed entity (although the identifier does not necessarily have any implied time ordering); 2.) a per team member unique identifier, which is ordered based on the time of creation of the partition of the

teamed entity created on that team member (a local team member database row identification, for example); and 3.) a creation timestamp for that partition of the teamed entity created on that team member. It is noted that the clocks on the team members 170 are synchronized, or time-aligned, within a tight enough tolerance to allow the timestamp to be used. For example, in accordance with some implementations, Network Time Protocol (NTP) synchronization of clients may be used.

[0060] The non-ordered team-wide unique identification allows the client application 134 to identify records, which match across team members 170, i.e., identify "stitch points." The ordered per team member unique identifier allows the client application 134 to retrieve the next/previous page of results from each team member 170 and therefore, implement a forward/reverse sliding window for each team member 170, which may be used in a multi-wave merge operation. The creation timestamp allows the client application 134 to decide how far the client application needs to search down each team member's results to find the stitch points.

[0061] As a more specific example, Fig. 10 depicts pages retrieved from team member 1 (via pages depicted in column 1010), team member 2 (via pages depicted in column 1012) and team member 3 (via pages depicted in column 1014). The results from the team members are separated at page boundaries 1030 and 1032. Fig. 10 also depicts a column 1020 of results. For the following example, the client application 134 retrieves a page of up to two team member objects that are timed from each team member 170 and returns a page two teamed objects that are timed to the user in the results column 1020.

[0062] More specifically, in order to return the first page of results (A, B), the client application 134 reads one page of results from team member 1, which contains the first two objects (by order of cross-team creation time): A and B; two pages of results from team member 2; and two pages of results from team member 3. The teamed objects B and C, for this example, were actually

created at the same time from two different clients; but due to timing differences, teamed objects B and C were created in different order on team member 1 versus team member 2. Because of this, an extra page of results is read from team member 2 for purposes of determining whether a record for teamed object B could be found. The client application 134 knew that there was a record for teamed object B, as the record team member 1 had the information in it as to which team members the teamed object was created on. Moreover, the client application 134 knew that the first page of return results from team member 2 were still around the teamed object was created, so the client application determined that it was realistic to load an extra page to find it.

[0063] For the first page of results, the results for team member 3 did not include a record for teamed object B. In order to return the second page of results (C, D), the client application reads one further page of results from team member 1, which contains the next two objects: C and D. Moreover, for this second page of results, no further pages are read from team member 2, if two to three pages are cached for each team member 170, as the information for objects C and D are available in the two pages already cached. From these cached results, the client application 134 knows that it cannot find a record for teamed objects C or D for team member 3.

[0064] In order to return the third page of results (E, F), the client application 134 reads one further page of results from team member 1, which contains the next two objects: E and F. The client application further reads one page of results from team member 2, which contains the next two objects: E and F. No further pages of results are retrieved for team member 3, as object E was in the first page (cached). Moreover, the client application 134 knows that it would not find a record for team object F from the creation information in the record for team member 1.

[0065] In order to return the third page of results (X, Y), the client application 134 reads the following, no further page results are retrieved from team

member 1 (i.e., the end has been reached); no further pages of results are retrieved for team member 2 (i.e., the end has been reached); and one further page of results is retrieved from team member 3, which contains the next two objects: X and Y.

[0066] For purposes of returning the fourth page of results (Z), the client application 134 reads the following. No further page of results from team member 1 (i.e., the end has been reached); no further pages of results from team member 2 (i.e., the end is reached); and one further page of results from team member 3, which contains the next object: Z.

[0067] In accordance with example implementations, a manifest is created and maintained for each teamed object. In general, the manifest, called an "object manifest," herein, describes the details of the data for a given teamed object stored among the team members. In this manner, the manifest allows the system to track and consolidate the distributed individual data items into one cohesive teamed object. In accordance with example implementations that are disclosed herein, the object manifest is distributed among the team members 170.

[0068] More specifically, referring to Fig. 11, in accordance with an example implementation, a technique 1100 includes storing (block 1102) deduplicated data for an object on a plurality of team members 170, or stores, and distributing (block 1104) a manifest, which describes the storage of the teamed object among the stores. In this manner, for each store, the technique 1100 includes storing (block 1106) data for the manifest, which describes the storage of the chunks on that store.

[0069] Distributing the object manifest among the team members 170, which is unlike a single master manifest, may help avoid a single point of failure. In other words, with a single manifest, the manifest may become lost or corrupted, which may render the teamed object useless, regardless of the state of the underlying data objects. However, by distributing the object manifest, each team member's object manifest (part of the overall distributed

object manifest) is entirely self-describing. In other words, each team member 170 has knowledge where its chunks of data fit within the larger teamed object. By storing distributed data in this way, overhead may be reduced, while robustness, redundancy and flexibility may be increased.

[0070] Referring to Fig. 12, in accordance with an example implementation, a master manifest 1200 is created by distributing member manifests 1240 (member manifests 1240-1, 1240-2, 1240-3 and 1240-4, being depicted in Fig. 12 as an example for four respective team members 170), which are stored on individual team members 170. Each member manifest 1240, in turn, includes entries, with each describing the chunks for the associated teamed objects stored on that team member. For example, for team member 1 for the example of Fig. 12, the member manifest 1240-1 contains multiple entries 1244 (entries 1244-1 and 1244-2, being depicted in Fig. 12, as examples), which describe the storage of corresponding chunks. Continuing the example, the team member nodes 2, 3 and 4 store corresponding entries 1246, 1248 and 1250, respectively. Collectively, the entries 1244, 1246, 1248 and 1250 form the entries 1220 of the master manifest 1210.

[0071] Thus, as depicted in Fig. 12, the master manifest 1210 includes various entries 1220 (entries 1220-1, 1220-2, 1220-3 and 1220-4, being depicted in Fig. 12 as specific examples), which correspond to the entries that are distributed across the team members.

[0072] In general, each entry (where "entry" refers to the entries stored on the team member or collected as part of the member manifest 1240) contains four fields of information: 1.) a first field that specifies where the associated chunk of data resides in the teamed object; 2.) a second field that specifies where the block of data resides in the member object; 3.) a third field indicating the size (in bytes, for example) of the chunk; and 4.) a fourth field that contains data identifying the specific team member on which the associated chunk is stored.

[0073] Thus, as depicted in Fig. 12, with the distributed master manifest 1200, each team member 170 contains a member manifest 1240, which only describes the chunks, which the member stores locally.

[0074] In contrast to a given entry 1220 of the master manifest 1210, the corresponding member manifest entry contains less information. For example, as compared to the master manifest entry 1220, a member manifest entry does not identify a node as all data stored on the team member has the same node identifier. Instead, a field is added to the team member's object store, describing which team members make up the overall team for the teamed object. This has the added benefit of allowing a team member to be able to contact any of the other team members to find out which team members store data for a given teamed object. Additionally, the member offset in the member manifest entry is not present. In this regard, team members only use teamed offsets, as it is up to the team member regarding how to store their data.

[0075] During write operations, each team member 170 records in its member manifest 1240 the data regions that it possesses and where the corresponding chunks reside. When reconstructing the team catalyst chunk for a particular read operation, the corollary to the bidding concept may be used. Another message may be added to the protocol so that the client application 134 may retrieve from each team member about the chunks of data stored for a given teamed object (offset and size, for example).

[0076] Thus, the approach disclosed herein federates out the master manifest for a given teamed object among the team members along with the user data, thereby obviating storage of the master manifest at a single location somewhere else. The federated approach may help use fewer protocol messages for read operations and, in accordance with example implementations, no additional messages for writes, as the manifest is tightly coupled with the data on each team member 170.

[0077] Moreover, the loss of an object's master manifest may result in the loss of the entire object, while the loss of an individual member manifest may result in only a partial loss of the object. Moreover, the approach described herein avoids adding redundancy, as redundancy may be relatively complicated. For example, a redundant master manifest would track where each redundant chunk is stored. Also, if the master manifest was stored in more than one place, then each manifest would be synchronized with each other. Considering the case where one of the manifest copies becomes "damaged," significant challenges may exist in determining, with certainty, which of the other copy(ies) is the "good" one. Should a master manifest be completely lost or damaged, there may be no way to reconstruct it. In addition, it may be challenging to add or remove team members from an existing distributed object.

[0078] One way to increase the redundancy of the stored data is to store each chunk in more than one team member 170. Using the bidding process, the client application 134 may choose to store the top two bids (as an example) instead of the top one. This would mean that every region may be stored more than once and always on more than one server 110, albeit to the detriment of overall dedupability. Should data be lost on one team member 170, the teamed object may still be reconstructed from the remaining team member objects. The previous level of redundancy for the object may be reinstated by reading back the manifests of the remaining team members, identifying regions, within sufficient redundancy and then writing the amount to a new team member object. Using this approach, redundancy may be achieved with relatively little associated system overhead.

[0079] For purposes of migration, an end user may desired to migrate a team member 170 object to a different node with the ultimate goal being to store the object on a different set of disks to free up space on the original node 150. With the distributed member manifest approach, an exact copy of the object on the team member 170 may be migrated from and stored on the new team member 170 that is the target of the migration. The next step is to update the

list of team members 170 participating in the storage of that distributed object to remove the old team member 170 and add the new team member 170.

[0080] An end user may want to add or remove a particular team member. These operations may be performed using mechanisms similar to the migration described above and by running one or more data write jobs from one or more team members; and furthermore, updating a list of participants among all of the nodes.

[0081] If the client application 134 stores user data regions in multiple team member chunks for redundancy, this provides an opportunity for the client application 134 to select which team member 170 from which the user data may be read based on each team member 170 returning server loading information in its response to the "which regions do you own for this extent" message.

[0082] In accordance with example implementations, the backup application 132 may control the replication of data from one team member 170 to another team member 170 over a relatively low bandwidth connection. In this manner, in accordance with example implementations, each team member 170 includes copy job engines 1310 and client applications 1320, as depicted in Fig. 13. The copy job engine 1310 on each team member 170 is constructed to access any chunk in the teamed object via its own instance of the client application 1320. This allows any of the team members 170 to perform copy operations to another team member 170, without the destination teamed store having the same number of team members 170 or any of the same team members as the origin team store, thereby providing replication inoperability across the product portfolio. Therefore, as depicted in Fig. 13, a given copy job engine 1310 on a given team member 170 may use a relatively low bandwidth connection 1350 to replicate a first teamed object 1370 (stored on team members 170-1, 170-2 and 170-3) to form a corresponding different team object 1372, which may have, as illustrated in Fig. 13, be stored on a

different grouping of team members 170 (i.e., two team members 170-4 and 170-5, for this example).

[0083] For purposes of avoiding rehydration the data during the replication copy, in accordance with some implementations, the client application 134 provides application programming interfaces (APIs) for non-hydrated read and write data paths. In this manner, a non-hydrated read includes reading hashes and unique chunk data with the client application internally dealing with acquiring region information from each team member for purposes of learning where to read the data from. A non-hydrated write operation includes matching hashes and storing unique chunks, with the client application 134 internally dealing the bidding and routing.

[0084] Fig. 14 illustrates a non-hydrated copy 1400 involving a source client application 1410, an associated copy job engine 1420 and a target catalyst application 1430. For each section of the copy (a ten megabyte (MB) selection, for example), the copy engine 1420 requests 1440 a manifest of hashes from the source client application 1400 and sends 1450 these hashes to a target 1460 to be matched. The target 1460 responds 1470 with a list of unmatched hashes for which chunk data are requested. The copy engine 1420 then requests 1480 these unique chunks from the source client application 1435, receivers 1486 the unmatched chunks and then sends 1488 them to the target 1460 to be stored.

[0085] While a limited number of examples have been disclosed herein, numerous modifications and variations therefrom can be appreciated. It is intended that the appended claims cover all such modifications and variations.

What is claimed is:

1 1. A method comprising:
2 communicating signatures of samples of first data associated with an object to
3 at least some stores of a plurality of stores on which the object is distributed;
4 in response to the communication, receiving responses indicating numbers of
5 the samples stored on the respective at least some of the plurality of stores; and
6 regulating on which store of the plurality of stores deduplicated data
7 associated with the first data is stored based at least in part on the numbers and a
8 pattern of data storage on the stores.

1 2. The method of claim 1, wherein the regulating comprises biasing
2 storage of the deduplicated data to be on one of the stores on which most recently
3 deduplicated data associated with the object was stored.

1 3. The method of claim 1, wherein the regulating comprises determining
2 whether the first data comprises seeding data.

1 4. The method of claim 3, wherein the regulating further comprises based
2 at least in part on whether the first data comprises seeding data, determining
3 whether the one of the numbers exceeds a re-routing threshold.

1 5. The method of claim 3, wherein the regulating further comprises based
2 at least in part on whether the first data comprises seeding data, determining
3 whether one of the numbers indicates a store that is current bid winner and an
4 immediately previous bid winner.

1 6. The method of claim 1, wherein the regulating comprises storing the
2 deduplicated data on a substantially randomly selected store based on whether the
3 deduplicated data is initially stored backup data for a system.

1 7. The method of claim 1, further comprising selectively suspending the
2 communication of the signatures based at least in part on a prediction of future
3 bidding activity.

1 8. The method of claim 1, further comprising:
2 basing the determination on at least one of the following: a number of times
3 the store having the largest number changes, the numbers provided by stores that
4 did not have the highest numbers, the highest number, and an amount of data
5 communicated to the store having the highest number.

1 9. A system comprising:
2 a backup application comprising a processor to designate first data
3 associated with an object to be backed up, the object being stored on a plurality of
4 stores in a distributed fashion; and
5 a client application to:
6 in response to the first data, communicate signatures of samples of the
7 first data to a plurality of stores on which the object is distributely stored;
8 receive responses from the plurality of stores indicating numbers of the
9 samples stored on the respective at least some of the plurality of stores; and
10 regulate on which store of the plurality of stores deduplicated data
11 associated with the first data is stored based at least in part on the numbers and a
12 pattern of data storage on the stores.

1 10. The system of claim 9, wherein the client application is to bias storage
2 of the deduplicated data to be on one of the stores on which most recently
3 deduplicated data associated with the object was stored.

1 11. The system of claim 9, wherein the client application is to regulate
2 which store of the plurality of stores the first data is stored based at least in part on
3 whether the first data comprises seeding data.

1 12. The system of claim 11, wherein the client application is to compare
2 the numbers to a bid threshold to determine whether the first data comprises seeding
3 data.

1 13. An article comprising a computer readable non-transitory storage
2 medium to stored instructions that when executed by the computer cause the
3 computer to:
4 communicate signatures of samples of first data associated with an object to
5 at least some stores of a plurality of stores on which the object is distributed;
6 in response to the communication, receive responses indicating numbers of
7 the samples stored on the respective at least some of the plurality of stores; and
8 regulate on which store of the plurality of stores deduplicated data associated
9 with the first data is stored based at least in part on the numbers and a pattern of
10 data storage on the stores.

11

1 14. The article of claim 13, the storage medium storing instructions that
2 when executed by the computer cause the computer to regulate which store of the
3 plurality of stores the first data is stored based at least in part on whether the first
4 data comprises seeding data.

1 15. The article of claim 14, the storage medium storing instructions that
2 when executed by the computer cause the computer to compare the numbers to a
3 bid threshold to determine whether the first data comprises seeding data.

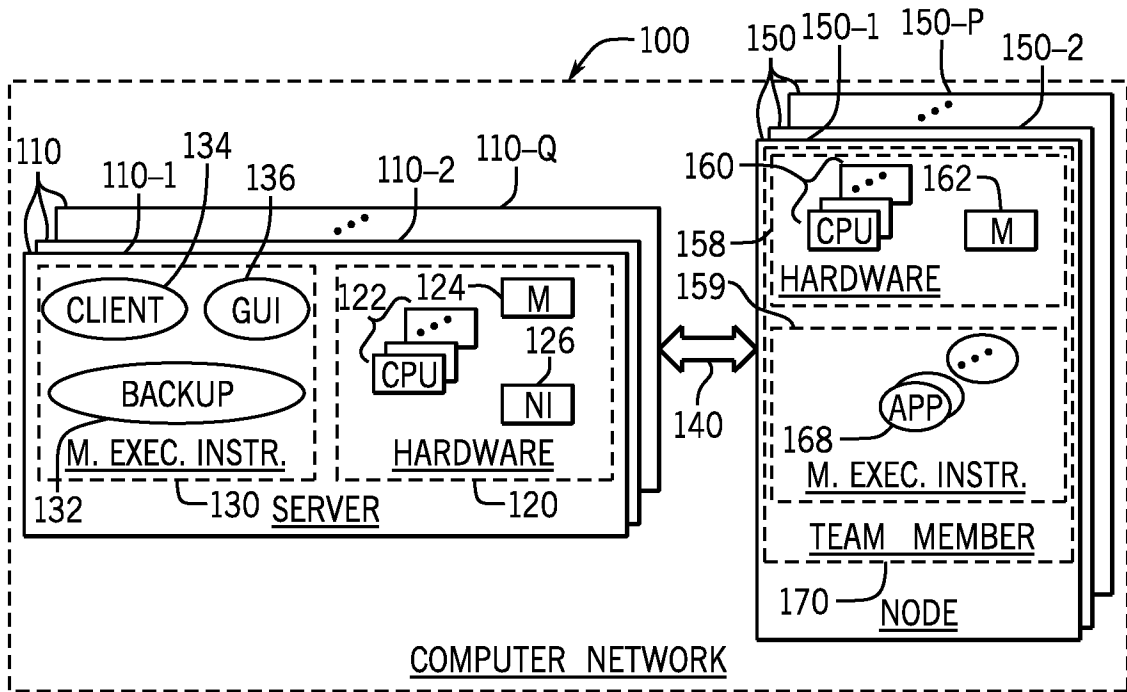


FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

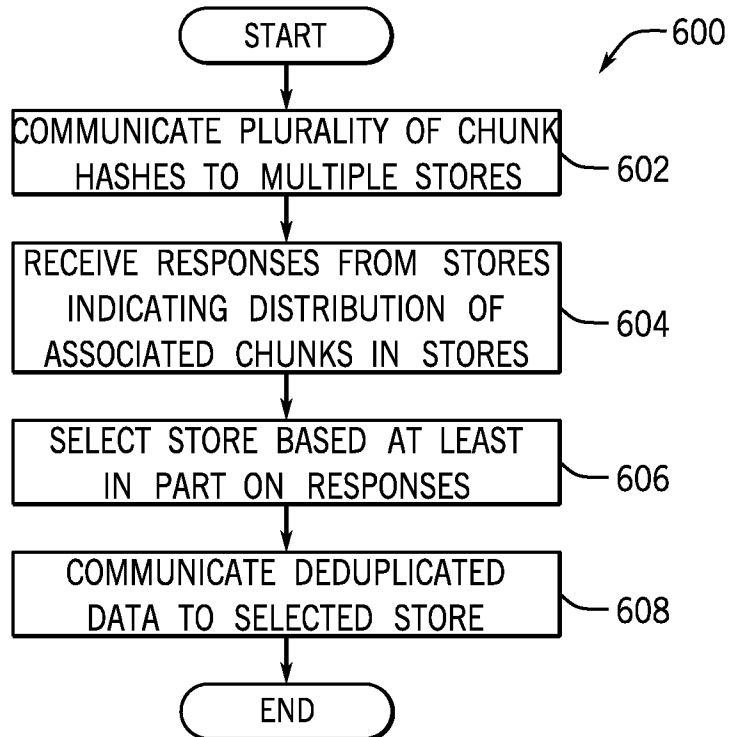


FIG. 6

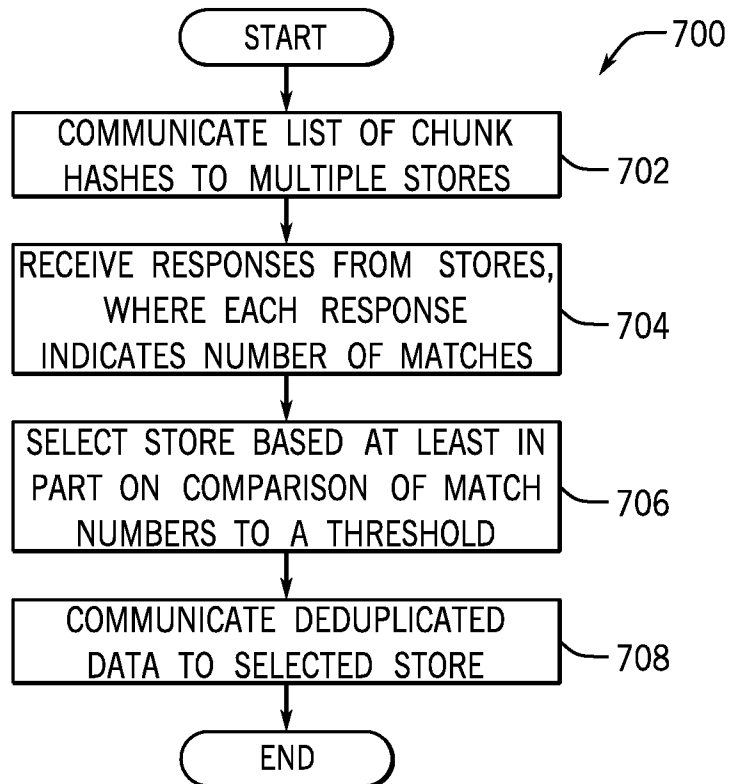


FIG. 7

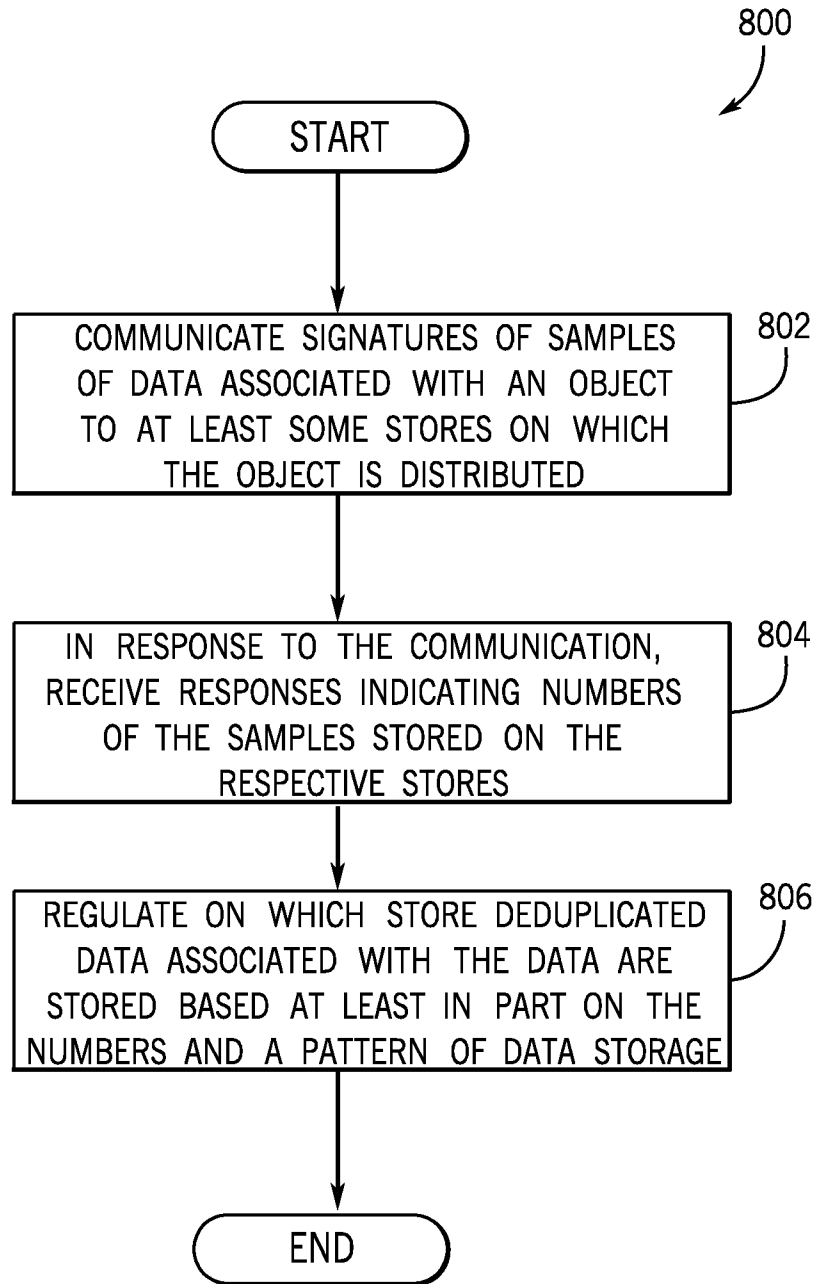


FIG. 8A

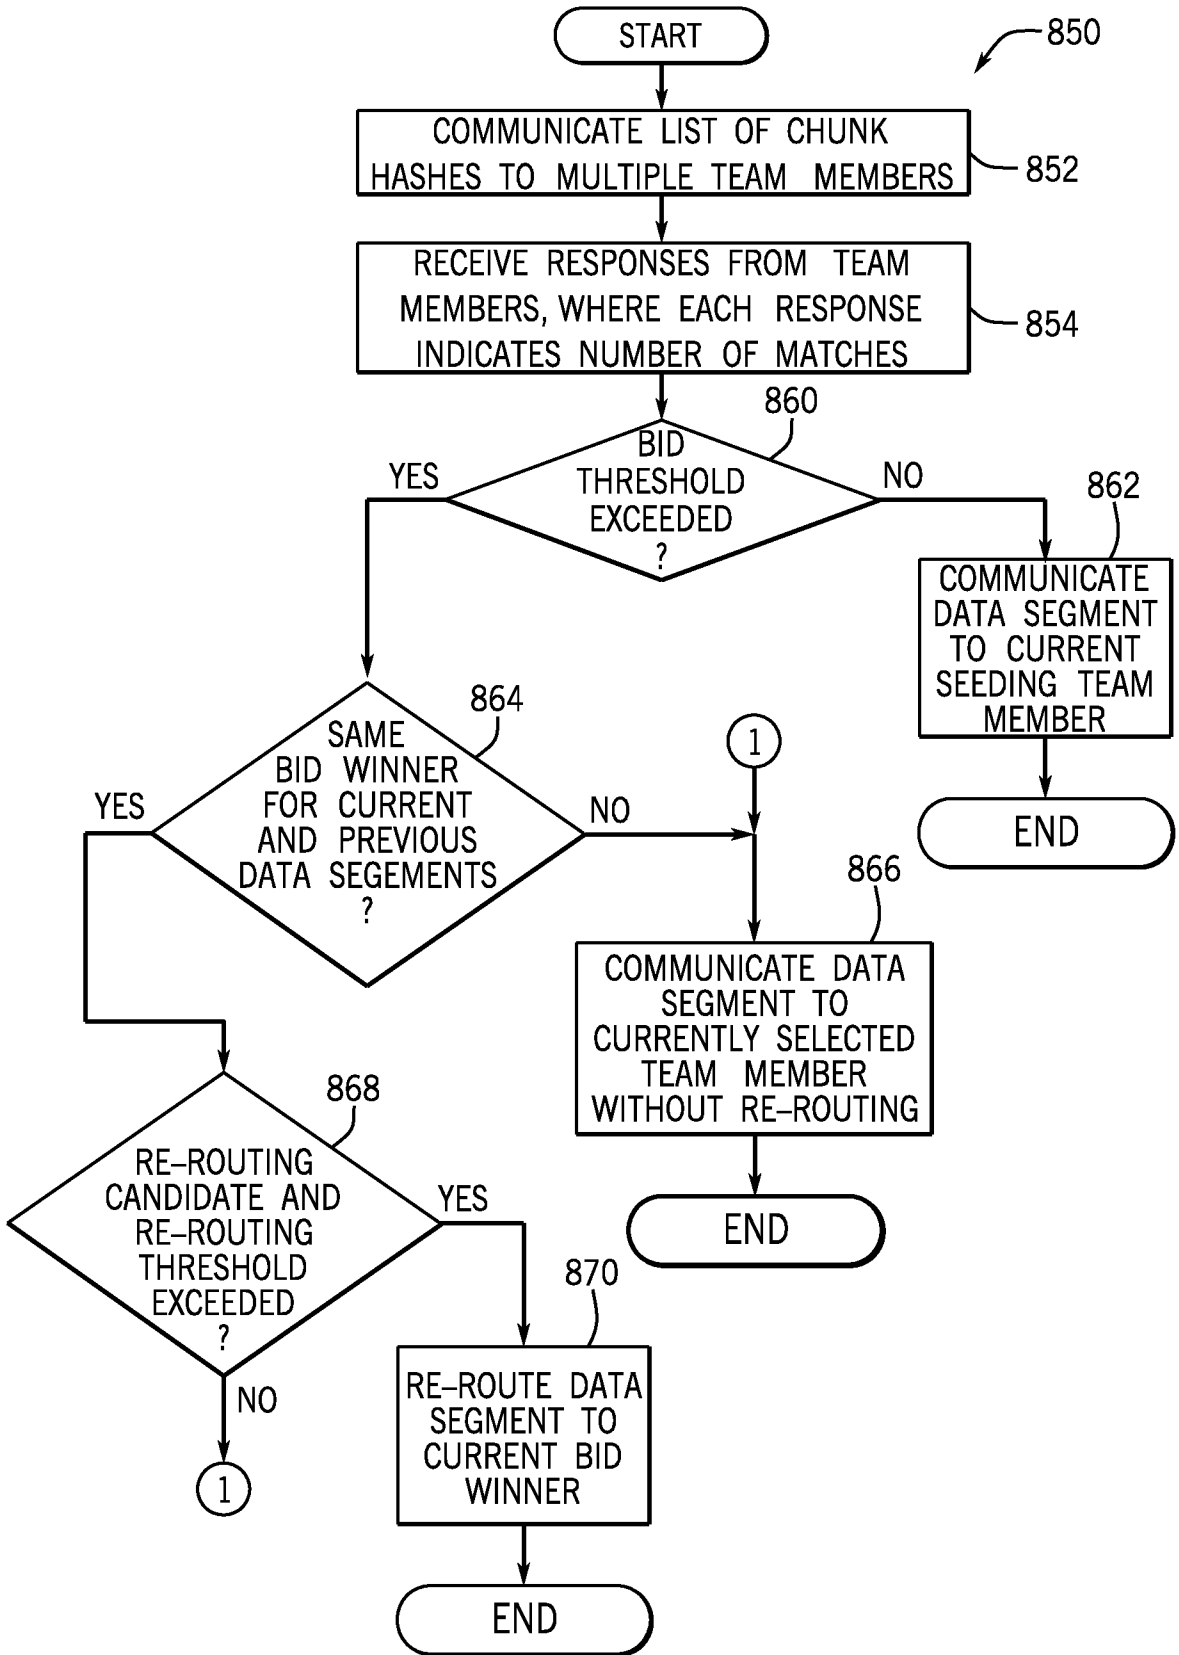


FIG. 8B

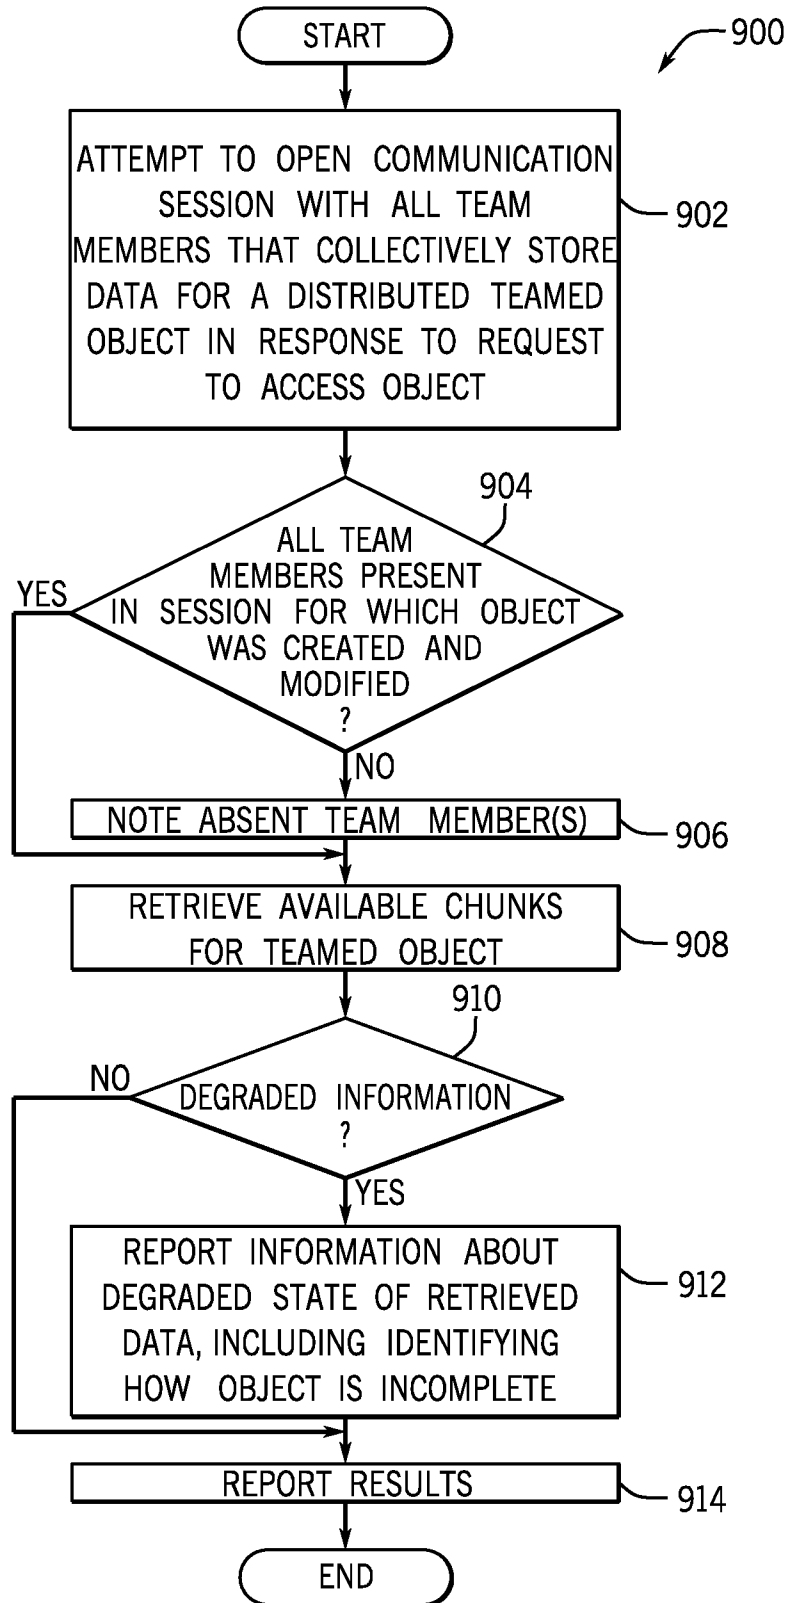


FIG. 9

1010 TEAM MEMBER 1			1012 TEAM MEMBER 2			1014 TEAM MEMBER 3			RESULTS 1020		
A			A			A			A		
B			C			E			A		
PAGE BOUNDARY											
C			B			X			B C		
D			D			Y			D		
PAGE BOUNDARY											
E			E			Z			E		
F			F			F			F X Y Z		

FIG. 10

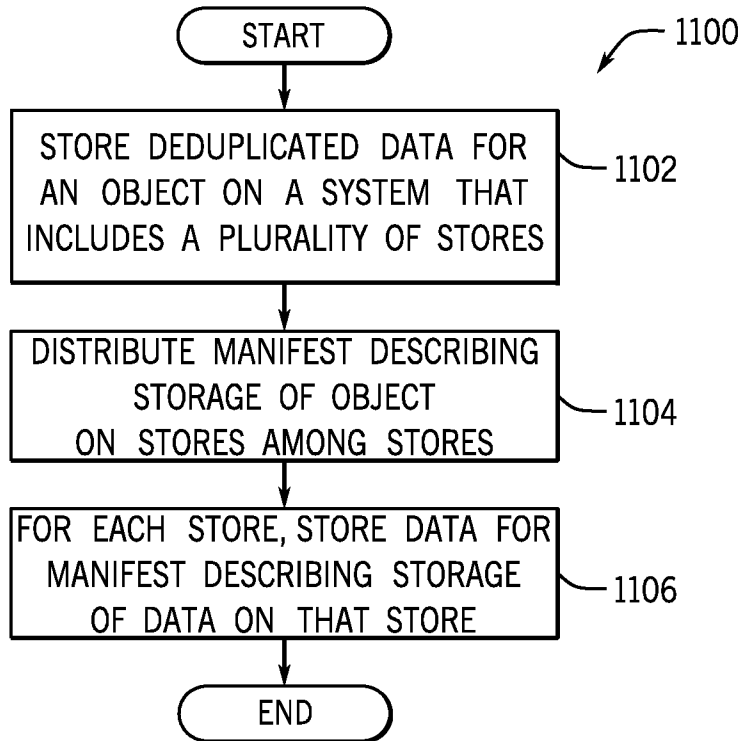


FIG. 11




FIG. 12




FIG. 13




FIG. 14

A. CLASSIFICATION OF SUBJECT MATTER**G06F 17/30(2006.01)i, G06F 12/14(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 17/30; G06F 12/02; G06F 12/16; G06F 12/14

Documentation 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 models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: deduplicate, file, data, object, signature, store, storage

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2013-0046944 A1 (SEIICHI DOMYO et al.) 21 February 2013 See paragraphs [0008]-[0010], [0032]-[0034], [0037]-[0041], [0073]-[0075]; and figures 1-4B.	1-15
A	US 2012-0084519 A1 (MANOJ KUMAR VIJAYAN et al.) 05 April 2012 See paragraphs [0005]-[0011], [0024]-[0038], [0050]-[0052]; and figures 1-2.	1-15
A	US 2012-0271798 A1 (MICHAEL HIRSCH) 25 October 2012 See paragraphs [0013]-[0028]; and figures 1, 4.	1-15
A	US 2012-0159098 A1 (CHUN HO CHEUNG et al.) 21 June 2012 See paragraphs [0052]-[0063]; and figure 1.	1-15
A	US 2012-0109894 A1 (GREGORY TAD KISHI) 03 May 2012 See paragraphs [0010]-[0026], [0043]-[0054]; and figure 1.	1-15

 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

18 February 2014 (18.02.2014)

Date of mailing of the international search report

20 February 2014 (20.02.2014)

Name and mailing address of the ISA/KR


 International Application Division
 Korean Intellectual Property Office
 189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
 Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

KIM, Seong Woo

Telephone No. +82-42-481-3348



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/041369

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013-0046944 A1	21/02/2013	US 8627026 B2 WO 2013-027231 A1	07/01/2014 28/02/2013
US 2012-0084519 A1	05/04/2012	US 2012-0084518 A1 US 8572340 B2 US 8578109 B2 WO 2012-044367 A1	05/04/2012 29/10/2013 05/11/2013 05/04/2012
US 2012-0271798 A1	25/10/2012	US 2012-158660 A1 US 8364641 B2 US 8458132 B2	21/06/2012 29/01/2013 04/06/2013
US 2012-0159098 A1	21/06/2012	CN 102567218 A WO 2012-083267 A2 WO 2012-083267 A3	11/07/2012 21/06/2012 27/12/2012
US 2012-0109894 A1	03/05/2012	CN 102308288 A EP 2353091 A1 US 2010-205389 A1 US 8108638 B2 US 8281099 B2 WO 2010-089187 A1	04/01/2012 10/08/2011 12/08/2010 31/01/2012 02/10/2012 12/08/2010