

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2018395933 B2**

(54) Title
Content management client synchronization service

(51) International Patent Classification(s)
G06F 16/178 (2019.01)

(21) Application No: **2018395933** (22) Date of Filing: **2018.12.17**

(87) WIPO No: **WO19/133321**

(30) Priority Data

(31) Number	(32) Date	(33) Country
15/868,518	2018.01.11	US
15/868,489	2018.01.11	US
15/868,511	2018.01.11	US
62/611,473	2017.12.28	US
15/868,505	2018.01.11	US

(43) Publication Date: **2019.07.04**

(44) Accepted Journal Date: **2021.04.01**

(71) Applicant(s)
Dropbox, Inc.

(72) Inventor(s)
GOLDBERG, Isaac;JAYAKAR, Sujay;LAI, John;YING, Robert;KOORAPATI, Nipunn;GUPTA, Gautam;SONG, Geoffry;JUBB, Elmer Charles

(74) Agent / Attorney
FPA Patent Attorneys Pty Ltd, ANZ Tower 161 Castlereagh Street, Sydney, NSW, 2000, AU

(56) Related Art
S UPPOOR et al: "Cloud Based Synchronization of Distributed File System Hierarchies", Cluster Computing Workshops and Posters, 20 September 2010

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization

International Bureau

(43) International Publication Date
04 July 2019 (04.07.2019)



(10) International Publication Number
WO 2019/133321 A1

(51) International Patent Classification:

G06F 16/178 (2019.01)

(21) International Application Number:

PCT/US2018/065940

(22) International Filing Date:

17 December 2018 (17.12.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/611,473	28 December 2017 (28.12.2017)	US
15/868,518	11 January 2018 (11.01.2018)	US
15/868,489	11 January 2018 (11.01.2018)	US
15/868,511	11 January 2018 (11.01.2018)	US
15/868,505	11 January 2018 (11.01.2018)	US

(71) Applicant: **DROPBOX, INC.** [US/US]; 333 Brannan Street, San Francisco, CA 94107 (US).

(72) Inventors: **GOLDBERG, Isaac**; Dropbox, Inc., 333 Brannan Street, San Francisco, CA 94107 (US). **JAYAKAR, Sujay**; Dropbox, Inc., 333 Brannan Street, San Francisco, CA 94107 (US). **LAI, John**; Dropbox, Inc., 333 Brannan Street, San Francisco, CA 94107 (US). **YING, Robert**; Dropbox, Inc., 333 Brannan Street, San Francisco, CA 94107 (US). **KOORAPATI, Nipunn**; Dropbox, Inc., 333 Brannan Street, San Francisco, CA 94107 (US). **GUPTA, Gautam**; Dropbox, Inc., 333 Brannan Street, San Francisco, CA 94107 (US). **SONG, Geoffry**; Dropbox, Inc., 333 Brannan Street, San Francisco, CA 94107 (US). **JUBB, Elmer Charles**; Dropbox, Inc., 333 Brannan Street, San Francisco, CA 94107 (US).

(74) Agent: **MCKNIGHT, Brian**; Polsinelli PC, 1401 I Street, NW, Suite 800, Washington, DC 20005 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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).

Published:

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

(54) Title: CONTENT MANAGEMENT CLIENT SYNCHRONIZATION SERVICE

(57) Abstract: The disclosed technology relates to a system configured to identify at least one difference between a sync tree and at least one of a remote tree and a local tree. The sync tree represents a known synced state between a server state and a file system state, the remote tree represents the server state, and the local tree represents the file system state. The system may further generate, based on the at least one difference, a set of operations configured to converge the server state and the file system state and manage the execution of the set of operations.



WO 2019/133321 A1

CONTENT MANAGEMENT CLIENT SYNCHRONIZATION SERVICE

CROSS-REFERENCE TO RELATED APPLICATIONS

[1] This application claims the priority to U.S. non-provisional application 15/868,489 filed on January 11, 2018, U.S. non-provisional application 15/868,518 filed on January 11, 2018, U.S. non-provisional application 15/868,511 filed on January 11, 2018, U.S. non-provisional application 15/868,505 filed on January 11, 2018 and to U.S. provisional application number 62/611,473, filed on December 28, 2017, all of which are expressly incorporated by reference herein in its entirety.

BACKGROUND

[2] Content management systems allow users to access and manage content items across multiple devices using a network. Some content management systems may allow users to share content items and provide additional features that aid users in collaborating using the content items. Content management systems generally store content items on servers and allow users access to the content items over a network. Some content management systems also allow for local copies to be stored on a client device in order to provide users with faster access to content items in a more natural interface (e.g., a native application or within the file system of the client device). Additionally, this allows the user to have access to the content items when the user is offline. Content management systems attempt to synchronize copies of a content item across a number of client devices and the servers so that each copy is identical. However, synchronization of content items is difficult and is associated with numerous technical obstacles.

SUMMARY

[2a] By way of clarification and for avoidance of doubt, as used herein and except where the context requires otherwise, the term "comprise" and variations of the term, such as "comprising", "comprises" and "comprised", are not intended to exclude further additions, components, integers or steps.

[2b] According to a first aspect of the invention there is provided a system comprising: a processor; and a computer-readable medium storing instructions that, when executed by

the processor, cause the processor to: generate a comparison between a sync tree and at least one of a remote tree and a local tree, wherein the remote tree represents a server state during a first time period and the local tree represents a file system state during the first time period, and wherein the sync tree represents a synced state between the server state and the file system state during a previous time period; identify, based on the comparison, at least one difference between a sync tree and at least one of a remote tree and a local tree; generate a set of operations based on the at least one difference, wherein the set of operations is configured to converge the server state and the file system state; and manage execution of the set of operations.

[2c] According to second aspect of the invention there is provided a computer readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to: compare a remote tree and a sync tree, wherein the remote tree represents a server state during a first time period, and wherein the sync tree represents a synced state between the server state and a file system state for content items stored on the computing system during a previous time period; determine, based on the comparison, at least one difference between the remote tree and the sync tree, the at least one difference indicating that the server state and the file system state are out of sync; generate, based on the determined at least one difference, a set of operations for the content items stored on the computing system, wherein the set of operations is configured to operate on the content items stored on the computing system to converge the server state and the file system state.

[2d] According to a third aspect of the invention there is provided a computer readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to: obtain a set of tree data structures including a remote tree representing a server state of content items associated with a user account on a content management system during a current time period, a local tree representing a file system state of content items associated with the user account on a client device during the current time period, and a sync tree representing a known sync state during a previous time period where the local tree and the remote tree were identical, wherein the sync tree matches the local tree and the remote tree at the previous time period; and between the content management system and the client device, wherein at the known sync state the

content items associated with the user account on the content management system are synchronized with the content items associated with the user account on the client device; and determine that the user account on the content management system includes at least one modification not synchronized to the client device by comparing the remote tree and the sync tree.

[2e] According to a fourth aspect of the invention there is provided a computer-implemented method for synchronizing modifications to a user account on a content management system to a client device authorized to access the user account, the method comprising: computing a difference between a remote tree data structure representing a server state for content items associated with the user account on the content management system and a sync tree data structure representing a known synchronization state between the content management system and the client device; and generating, based on the difference, a set of operations that when performed on the client device update the content items stored on the client device to converge a file system state on the client device and the server state.

[2f] According to a fifth aspect of the invention there is provided a computer-implemented method for synchronizing modifications to content items on a client device to a user account on a content management system, the method comprising: computing a difference between a local tree data structure representing a file system state for content items associated with the user account on the client device and a sync tree data structure representing a known synchronization state between the content management system and the client device; and generating, based on the difference, a set of operations that when performed update the content items stored on the content management system to converge a server state for content items associated the user account on the content management system and the file system state.

BRIEF DESCRIPTION OF THE DRAWINGS

[3] The above-recited and other advantages and features of the present technology will become apparent by reference to specific implementations illustrated in the appended drawings. A person of ordinary skill in the art will understand that these drawings only show some examples of the present technology and would not limit the scope of the present

2018395933 15 Feb 2021

technology to these examples. Furthermore, the skilled artisan will appreciate the principles of the present technology as described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[4] FIG. 1 shows an example of a content management system and client devices, in accordance with some embodiments;

[5] FIG. 2 shows an example of a client synchronization service, in accordance with some embodiments;

[6] FIG. 3 shows an example of tree data structures, in accordance with various embodiments;

[7] FIG. 4 shows an example of tree data structures, in accordance with various embodiments;

[8] FIG. 5 shows an example method for synchronizing a server state and a file system state using tree data structures, in accordance with various embodiments of the subject technology;

[9] FIG. 6 shows an example method for resolving conflicts when synchronizing a server state and a file system state using tree data structures, in accordance with various embodiments of the subject technology;

[10] FIG. 7 shows an example of tree data structures illustrating a violation of a rule for an add operation, in accordance with various embodiments;

[11] FIG. 8 shows an example method for incrementally converging a server state and a file system state, in accordance with various embodiments of the subject technology;

[12] FIG. 9 shows an example of tree data structures, in accordance with various embodiments;

[13] FIG. 10 shows an example scenario;

[14] FIG. 11 shows an example Venn diagram representation of two plans of operations, in accordance with various embodiments of the subject technology;

[15] FIG. 12 shows an example method for managing changes in plans of operations, in accordance with various embodiments of the subject technology;

[16] FIG. 13 shows an illustration of a filename array and a hash index array, in accordance with various embodiments of the subject technology;

[17] FIG. 14 shows an example method for storing a filename, in accordance with various embodiments of the subject technology;

[18] FIG. 15 shows an example method for retrieving a location of a filename given the filename, in accordance with various embodiments of the subject technology;

[19] FIGS. 16A and 16B show examples of tree data structures, in accordance with various embodiments;

[20] FIG. 17 shows an example of tree data structures, in accordance with various embodiments;

[21] FIG. 18 shows an example method for retrieving a location of a filename given the filename, in accordance with various embodiments of the subject technology; and

[22] FIG. 19 shows an example of a system for implementing certain aspects of the present technology.

DETAILED DESCRIPTION

[23] Various examples of the present technology are discussed in detail below. While specific implementations are discussed, it should be understood that this is done for illustration purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without parting from the spirit and scope of the present technology.

[24] Various advances in computing and networking technologies have enabled content management systems to provide users with access to content items across multiple devices. The content items may include, but are not limited to, files, documents, messages (e.g., email messages or text messages), media files (e.g., photos, videos, and audio files), folders containing other content items, or any other unit of content. Content items may be shared with multiple users, edited, deleted, added, renamed, or moved. However, synchronizing these content items across several computing devices (e.g., servers and client devices) and across several user accounts has remained flawed and rife with technological obstacles.

[25] To illustrate some of the technical obstacles, a first machine (e.g., a client device or server) may send communications to a second machine that provides information about how a user has modified content items managed by the content management system. These communications may be used by the second machine to

synchronize the content items on the second machine such that actions performed on content items on the first machine are reflected in content items on the second machine and the content items on the first machine are substantially identical to the content items on the second machine.

[26] However, there may be several communications sent and the communications may be received out of order as a result of various network routing protocols used by the one or more networks used to transmit the communications, the technical operations of the first or second machine, or some other reason. Furthermore, a user may be performing a large number of modifications to a large number of content items, undo previous modifications in a short amount of time, or quickly perform additional modifications to a previously modified content item or set of content items. This increases the likelihood that these communications are received out of order, certain communications are out of date, or that the second machine will perform operations on content items that are not up to date. As a result, many of the operations may not be compatible with the current state of the content items. In fact, it may be difficult to even detect whether some operations are in conflict with other operations or with the current state of the content items.

[27] Additionally, there is an inherent latency with respect to synchronization actions. For example, actions taken on the first machine are first detected by the first machine, and a communication is generated and then transmitted through a network. The communication is received by the second machine, which may still be processing previous communications and taking actions detailed in the communications. In this illustrative scenario, there are several points where latency is introduced by limited computing resources (e.g., bandwidth, memory, processing time, processing cycles, etc.) of the first machine, the second machine, and/or the network. As latency increases the likelihood that communications, for some reason, conflict with the current state of the content items are increased. Furthermore, processing these conflicted communications and resolving the conflicts also expends needless computing resources such as processing time, memory, energy, or bandwidth and further increases latency.

[28] To further complicate matters, the same or different user on the second machine and/or additional machines with access to the content items may also be performing modification to the content items. As a result, the issues above may be

multiplied and additional technical issues arise as to whether local actions conflict with remote actions and/or whether local actions are operating on up to date content items.

[29] The disclosed technology addresses the need in the art for a client synchronization service for a content management system that provides a technical solution to the technical problems above as well as others. The client synchronization service may be configured to operate on a client device and identify synchronization mismatches between content items on a server of the content management system and corresponding content items on the client device. For each synchronization mismatch, the client synchronization service may identify operations needed to synchronize the content items and initiate those operations.

[30] The client synchronization service may track the status of content items on the server, the status of content items on the client device, and their synchronization state using a set of tree data structures (“trees”). According to some embodiments, a set of 3 trees may be used. The three trees may include a remote tree that represents a server state, a local tree that represents the file system state on the client device, and a sync tree that represents a merge base for the local tree and the remote tree. The merge base may be thought of as a common ancestor of the local tree and the remote tree or a last known synced state between the local tree and the remote tree. Accordingly, the client synchronization service may determine that the server state and the client device state are synchronized when all 3 trees (e.g., the remote tree, the sync tree, and the local tree) are identical.

[31] When a modification to the server state of the content items or the client device file system state (“file system state”) of the content items is detected, the client synchronization service updates the appropriate tree and determines whether the server state and the file system state are synchronized based on the triumvirate of trees. Based on the update to one of the trees, the server state and the file system state may become synchronized, become unsynchronized, or become further unsynchronized. If the server state and the file system state are not synchronized, the client synchronization service may identify at least an initial set of operations needed to converge the server state and the file system state and get the server state and the file system state closer to a synchronized state.

[32] By relying on the set of tree data structures to monitor the server state and the file system state provides alternatives and/or solutions rooted in computing technology to various technical problems. For example, the client synchronization service is able to track the server state as well as the file state and store a representation of a merge base of the two states. As a result, the various embodiments of the subject technology avoid the technical problems associated with receiving a number of communications specifying how users are modifying content items remotely and determining which order these modifications should be implemented locally, whether the modifications conflict with other modifications or are out of date, and whether remote modifications conflict with local modifications performed locally by users. Many of these issues arise from other solutions not being able to track the state of the various actors involved (e.g., the server and the client device) and not being able to quickly determine whether the states are in sync. Instead, these other solutions rely on receiving instructions on how to modify content items locally, without the context of whether the server state and file system state are in sync.

[33] Furthermore, since the server state and the file system state are continuously monitored, determining whether they are synced is much more efficient in terms of procedural complexity as well as computing time and resources. As is described in further detail below, the client synchronization service enables the incremental and methodical synchronization of the server state and the file system state in a more deterministic manner. As a result, the scaling and testing of content management system features is also more efficient.

Content Management System

[34] In some embodiments, the disclosed technology is deployed in the context of a content management system having content item synchronization capabilities and collaboration features, among others. An example system configuration 100 is shown in FIG. 1A, which depicts content management system 110 interacting with client device 150.

[35] Accounts

[36] Content management system 110 can store content items in association with accounts, as well as perform a variety of content item management tasks, such as retrieve, modify, browse, and/or share the content item(s). Furthermore, content

management system 110 can enable an account to access content item(s) from multiple client devices.

[37] Content management system 110 supports a plurality of accounts. An entity (user, group of users, team, company, etc.) can create an account with content management system, and account details can be stored in account database 140. Account database 140 can store profile information for registered entities. In some cases, profile information for registered entities includes a username and/or email address. Account database 140 can include account management information, such as account type (e.g. various tiers of free or paid accounts), storage space allocated, storage space used, client devices 150 having a registered content management client application 152 resident thereon, security settings, personal configuration settings, etc.

[38] Account database 140 can store groups of accounts associated with an entity. Groups can have permissions based on group policies and/or access control lists, and members of the groups can inherit the permissions. For example, a marketing group can have access to one set of content items while an engineering group can have access to another set of content items. An administrator group can modify groups, modify user accounts, etc.

[39] Content Item Storage

[40] A feature of content management system 110 is the storage of content items, which can be stored in content storage 142. Content items can be any digital data such as documents, collaboration content items, text files, audio files, image files, video files, webpages, executable files, binary files, etc. A content item can also include collections or other mechanisms for grouping content items together with different behaviors, such as folders, zip files, playlists, albums, etc. A collection can refer to a folder, or a plurality of content items that are related or grouped by a common attribute. In some embodiments, content storage 142 is combined with other types of storage or databases to handle specific functions. Content storage 142 can store content items, while metadata regarding the content items can be stored in metadata database 146. Likewise, data regarding where a content item is stored in content storage 142 can be stored in content directory 144. Additionally, data regarding changes, access, etc. can be stored in server file journal 148. Each of the various storages/databases such as content storage 142, content directory 144, server

file journal 148, and metadata database 146 can be comprised of more than one such storage or database and can be distributed over many devices and locations. Other configurations are also possible. For example, data from content storage 142, content directory 144, server file journal 148, and/or metadata database 146 may be combined into one or more content storages or databases or further segmented into additional content storages or databases. Thus, content management system 110 may include more or less storages and/or databases than shown in FIG. 1.

[41] In some embodiments, content storage 142 is associated with at least one content storage service 116, which includes software or other processor executable instructions for managing the storage of content items including, but not limited to, receiving content items for storage, preparing content items for storage, selecting a storage location for the content item, retrieving content items from storage, etc. In some embodiments, content storage service 116 can divide a content item into smaller chunks for storage at content storage 142. The location of each chunk making up a content item can be recorded in content directory 144. Content directory 144 can include a content entry for each content item stored in content storage 142. The content entry can be associated with a unique ID, which identifies a content item.

[42] In some embodiments, the unique ID, which identifies a content item in content directory 144, can be derived from a deterministic hash function. This method of deriving a unique ID for a content item can ensure that content item duplicates are recognized as such since the deterministic hash function will output the same identifier for every copy of the same content item, but will output a different identifier for a different content item. Using this methodology, content storage service 116 can output a unique ID for each content item.

[43] Content storage service 116 can also designate or record a content path for a content item in metadata database 146. The content path can include the name of the content item and/or folder hierarchy associated with the content item. For example, the content path can include a folder or path of folders in which the content item is stored in a local file system on a client device. While content items are stored in content storage 142 in blocks and may not be stored under a tree like directory structure, such directory structure is a comfortable navigation structure for users. Content storage service 116 can define or record a content path for a content item wherein the “root” node of a directory structure can be a namespace for each account.

Within the namespace can be a directory structure defined by a user of an account and/or content storage service 116. Metadata database 146 can store the content path for each content item as part of a content entry.

[44] In some embodiments the namespace can include additional namespaces nested in the directory structure as if they are stored within the root node. This can occur when an account has access to a shared collection. Shared collections can be assigned their own namespace within content management system 110. While some shared collections are actually a root node for the shared collection, they are located subordinate to the account namespace in the directory structure, and can appear as a folder within a folder for the account. As addressed above, the directory structure is merely a comfortable navigation structure for users, but does not correlate to storage locations of content items in content storage 142.

[45] While the directory structure in which an account views content items does not correlate to storage locations at content management system 110, the directory structure can correlate to storage locations on client device 150 depending on the file system used by client device 150.

[46] As addressed above, a content entry in content directory 144 can also include the location of each chunk making up a content item. More specifically, the content entry can include content pointers that identify the location in content storage 142 of the chunks that make up the content item.

[47] In addition to a content path and content pointer, a content entry in content directory 144 can also include a user account identifier that identifies the user account that has access to the content item and/or a group identifier that identifies a group with access to the content item and/or a namespace to which the content entry belongs.

[48] Content storage service 116 can decrease the amount of storage space required by identifying duplicate content items or duplicate blocks that make up a content item or versions of a content item. Instead of storing multiple copies, content storage 142 can store a single copy of the content item or block of the content item and content directory 144 can include a pointer or other mechanism to link the duplicates to the single copy.

[49] Content storage service 116 can also store metadata describing content items, content item types, folders, file path, and/or the relationship of content items to various accounts, collections, or groups in metadata database 146, in association with the unique ID of the content item.

[50] Content storage service 116 can also store a log of data regarding changes, access, etc. in server file journal 148. Server file journal 148 can include the unique ID of the content item and a description of the change or access action along with a time stamp or version number and any other relevant data. Server file journal 148 can also include pointers to blocks affected by the change or content item access. Content storage service can provide the ability to undo operations, by using a content item version control that tracks changes to content items, different versions of content items (including diverging version trees), and a change history that can be acquired from the server file journal 148.

[51] Content Item Synchronization

[52] Another feature of content management system 110 is synchronization of content items with at least one client device 150. Client device(s) can take different forms and have different capabilities. For example, client device 150₁ is a computing device having a local file system accessible by multiple applications resident thereon. Client device 150₂ is a computing device wherein content items are only accessible to a specific application or by permission given by the specific application, and the content items are typically stored either in an application specific space or in the cloud. Client device 150₃ is any client device accessing content management system 110 via a web browser and accessing content items via a web interface. While example client devices 150₁, 150₂, and 150₃ are depicted in form factors such as a laptop, mobile device, or web browser, it should be understood that the descriptions thereof are not limited to devices of these example form factors. For example a mobile device such as client 150₂ might have a local file system accessible by multiple applications resident thereon, or client 150₂ might access content management system 110 via a web browser. As such, the form factor should not be considered limiting when considering client 150's capabilities. One or more functions described herein with respect to client device 150 may or may not be available on every client device depending on the specific capabilities of the device – the file access model being one such capability.

[53] In many embodiments, client devices are associated with an account of content management system 110, but in some embodiments client devices can access content using shared links and do not require an account.

[54] As noted above, some client devices can access content management system 110 using a web browser. However, client devices can also access content management system 110 using client application 152 stored and running on client device 150. Client application 152 can include a client synchronization service 156.

[55] Client synchronization service 156 can be in communication with server synchronization service 112 to synchronize changes to content items between client device 150 and content management system 110.

[56] Client device 150 can synchronize content with content management system 110 via client synchronization service 156. The synchronization can be platform agnostic. That is, content can be synchronized across multiple client devices of varying type, capabilities, operating systems, etc. Client synchronization service 156 can synchronize any changes (new, deleted, modified, copied, or moved content items) to content items in a designated location of a file system of client device 150.

[57] Content items can be synchronized from client device 150 to content management system 110, and vice versa. In embodiments wherein synchronization is from client device 150 to content management system 110, a user can manipulate content items directly from the file system of client device 150, while client synchronization service 156 can monitor directory on client device 150 for changes to files within the monitored folders.

[58] When client synchronization service 156 detects a write, move, copy, or delete of content in a directory that it monitors, client synchronization service 156 can synchronize the changes to content management system service 116. In some embodiments, client synchronization service 156 can perform some functions of content management system service 116 including functions addressed above such as dividing the content item into blocks, hashing the content item to generate a unique identifier, etc. Client synchronization service 156 can index content within client storage index 164 and save the result in storage index 164. Indexing can include storing paths plus a unique server identifier, and a unique client identifier for each content item. In some embodiments, client synchronization service 156 learns the

unique server identifier from server synchronization service 112, and learns the unique client identifier from the operating system of client device 150.

[59] Client synchronization service 156 can use storage index 164 to facilitate the synchronization of at least a portion of the content within client storage with content associated with a user account on content management system 110. For example, client synchronization service 156 can compare storage index 164 with content management system 110 and detect differences between content on client storage and content associated with a user account on content management system 110. Client synchronization service 156 can then attempt to reconcile differences by uploading, downloading, modifying, and deleting content on client storage as appropriate. Content storage service 116 can store the changed or new block for the content item and update server file journal 148, metadata database 146, content directory 144, content storage 142, account database 140, etc., as appropriate.

[60] When synchronizing from content management system 110 to client device 150, a mount, modification, addition, deletion, move of a content item recorded in server file journal 148 can trigger a notification to be sent to client device 150 using notification service 117. When client device 150 is informed of the change a request changes listed in server file journal 148 since the last synchronization point known to the client device. When client device 150 determines that it is out of synchronization with content management system 110, client synchronization service 156 requests content item blocks including the changes, and updates its local copy of the changed content items.

[61] In some embodiments, storage index 164 stores tree data structures wherein one tree reflects the latest representation of a directory according to server synchronization service 112, while another tree reflects the latest representation of the directory according to client synchronization service 156. Client synchronization service can work to ensure that the tree structures match by requesting data from server synchronization service 112 or committing changes on client device 150 to content management system 110.

[62] Sometimes client device 150 might not have a network connection available. In this scenario, client synchronization service 156 can monitor the linked collection for content item changes and queue those changes for later synchronization to content

management system 110 when a network connection is available. Similarly, a user can manually start, stop, pause, or resume synchronization with content management system 110.

[63] Client synchronization service 156 can synchronize all content associated with a particular user account on content management system 110. Alternatively, client synchronization service 156 can selectively synchronize a portion of the content of the total content associated with the particular user account on content management system 110. Selectively synchronizing only a portion of the content can preserve space on client device 150 and save bandwidth.

[64] In some embodiments, client synchronization service 156 selectively stores a portion of the content associated with the particular user account and stores placeholder content items in client storage for the remainder portion of the content. For example, client synchronization service 156 can store a placeholder content item that has the same filename, path, extension, metadata, of its respective complete content item on content management system 110, but lacking the data of the complete content item. The placeholder content item can be a few bytes or less in size while the respective complete content item might be significantly larger. After client device 150 attempts to access the content item, client synchronization service 156 can retrieve the data of the content item from content management system 110 and provide the complete content item to accessing client device 150. This approach can provide significant space and bandwidth savings while still providing full access to a user's content on content management system 110.

[65] Collaboration features

[66] Another feature of content management system 110 is to facilitate collaboration between users. Collaboration features include content item sharing, commenting on content items, co-working on content items, instant messaging, providing presence and seen state information regarding content items, etc.

[67] Sharing

[68] Content management system 110 can manage sharing content via sharing service 128. Sharing content by providing a link to the content can include making the content item accessible from any computing device in network communication with content management system 110. However, in some embodiments a link can be

associated with access restrictions enforced by content management system 110 and access control list 145. Sharing content can also include linking content using sharing service 128 to share content within content management system 110 with at least one additional user account (in addition to the original user account associated with the content item) so that each user account has access to the content item. The additional user account can gain access to the content by accepting the content, which will then be accessible through either web interface service 124 or directly from within the directory structure associated with their account on client device 150. The sharing can be performed in a platform agnostic manner. That is, the content can be shared across multiple client devices 150 of varying type, capabilities, operating systems, etc. The content can also be shared across varying types of user accounts.

[69] To share a content item within content management system 110 sharing service 128 can add a user account identifier or multiple user account identifiers to a content entry in access control list database 145 associated with the content item, thus granting the added user account access to the content item. Sharing service 128 can also remove user account identifiers from a content entry to restrict a user account's access to the content item. Sharing service 128 can record content item identifiers, user account identifiers given access to a content item, and access levels in access control list database 145. For example, in some embodiments, user account identifiers associated with a single content entry can specify different permissions for respective user account identifiers with respect to the associated content item.

[70] To share content items outside of content management system 110, sharing service 128 can generate a custom network address, such as a uniform resource locator (URL), which allows any web browser to access the content item or collection in content management system 110 without any authentication. To accomplish this, sharing service 128 can include content identification data in the generated URL, which can later be used to properly identify and return the requested content item. For example, sharing service 128 can include the account identifier and the content path or a content item identifying code in the generated URL. Upon selection of the URL, the content identification data included in the URL can be transmitted to content management system 110, which can use the received content identification data to identify the appropriate content item and return the content item.

[71] In addition to generating the URL, sharing service 128 can also be configured to record in access control list database 145 that a URL to the content item has been created. In some embodiments, the content entry associated with a content item can include a URL flag indicating whether a URL to the content item has been created. For example, the URL flag can be a Boolean value initially set to 0 or false to indicate that a URL to the content item has not been created. Sharing service 128 can change the value of the flag to 1 or true after generating a URL to the content item.

[72] In some embodiments, sharing service 128 can associate a set of permissions to a URL for a content item. For example, if a user attempts to access the content item via the URL, sharing service 128 can provide a limited set of permissions for the content item. Examples of limited permissions include restrictions that the user cannot download the content item, save the content item, copy the content item, modify the content item, etc. In some embodiments, limited permissions include restrictions that only permit a content item to be accessed from with a specified domain, i.e., from within a corporate network domain, or by accounts associated with a specified domain, e.g., accounts associated with a company account (e.g., @acme.com).

[73] In some embodiments, sharing service 128 can also be configured to deactivate a generated URL. For example, each content entry can also include a URL active flag indicating whether the content should be returned in response to a request from the generated URL. For example, sharing service 128 can only return a content item requested by a generated link if the URL active flag is set to 1 or true. Thus, access to a content item for which a URL has been generated can be easily restricted by changing the value of the URL active flag. This allows a user to restrict access to the shared content item without having to move the content item or delete the generated URL. Likewise, sharing service 128 can reactivate the URL by again changing the value of the URL active flag to 1 or true. A user can thus easily restore access to the content item without the need to generate a new URL.

[74] In some embodiments, content management system 110 can designate a URL for uploading a content item. For example, a first user with a user account can request such a URL, provide the URL to a contributing user and the contributing user can upload a content item to the first user's user account using the URL.

[75] Team Service

[76] In some embodiments content management system 110 includes team service 130. Team service 130 can provide functionality for creating and managing defined teams of user accounts. Teams can be created for a company, with sub-teams (e.g., business units, or project teams, etc.), and user accounts assigned to teams and sub-teams, or teams can be created for any defined group of user accounts. Team's service 130 can provide a common shared space for the team, private user account folders, and access limited shared folders. Team's service can also provide a management interface for an administrator to manage collections and content items within team, and can manage user accounts that are associated with the team.

[77] Authorization Service

[78] In some embodiments, content management system 110 includes authorization service 132. Authorization service 132 ensures that a user account attempting to access a namespace has appropriate rights to access the namespace. Authorization service 132 can receive a token from client application 152 that follows a request to access a namespace and can return the capabilities permitted to the user account. For user accounts with multiple levels of access (e.g. a user account with user rights and administrator rights) authorization service 132 can also require explicit privilege escalation to avoid unintentional actions by administrators.

[79] Presence and Seen State

[80] In some embodiments, content management system can provide information about how users with which a content item is shared are interacting or have interacted with the content item. In some embodiments, content management system 110 can report that a user with which a content item is shared is currently viewing the content item. For example, client collaboration service 160 can notify notifications service 117 when client device 150 is accessing the content item. Notifications service 117 can then notify all client devices of other users having access to the same content item of the presence of the user of client device 150 with respect to the content item.

[81] In some embodiments, content management system 110 can report a history of user interaction with a shared content item. Collaboration service 126 can query data sources such as metadata database 146 and server file journal 148 to determine that a user has saved the content item, that a user has yet to view the content item, etc., and

disseminate this status information using notification service 117 to other users so that they can know who currently is or has viewed or modified the content item.

[82] Collaboration service 126 can facilitate comments associated with content, even if a content item does not natively support commenting functionality. Such comments can be stored in metadata database 146.

[83] Collaboration service 126 can originate and transmit notifications for users. For example, a user can mention another user in a comment and collaboration service 126 can send a notification to that user that he has been mentioned in the comment. Various other content item events can trigger notifications, including deleting a content item, sharing a content item, etc.

[84] Collaboration service 126 can provide a messaging platform whereby users can send and receive instant messages, voice calls, emails, etc.

[85] Collaboration Content Items

[86] In some embodiments content management service can also include Collaborative document service 134 which can provide an interactive content item collaboration platform whereby users can simultaneously create collaboration content items, comment in the collaboration content items, and manage tasks within the collaboration content items. Collaboration content items can be files that users can create and edit using a collaboration content item editor, and can contain collaboration content item elements. Collaboration content item elements may include a collaboration content item identifier, one or more author identifiers, collaboration content item text, collaboration content item attributes, interaction information, comments, sharing users, etc. Collaboration content item elements can be stored as database entities, which allows for searching and retrieving the collaboration content items. Multiple users may access, view, edit, and collaborate on collaboration content items at the same time or at different times. In some embodiments this can be managed by requiring two users access a content item through a web interface and there they can work on the same copy of the content item at the same time.

[87] Collaboration Companion Interface

[88] In some embodiments client collaboration service 160 can provide a native application companion interface for the purpose of displaying information relevant to a content item being presented on client device 150. In embodiments wherein a

content item is accessed by a native application stored and executed on client device 150, where the content item is in a designated location of the file system of client device 150 such that the content item is managed by content application 152, the native application may not provide any native way to display the above addressed collaboration data. In such embodiments, client collaboration service 160 can detect that a user has opened a content item, and can provide an overlay with additional information for the content item, such as collaboration data. For example, the additional information can include comments for the content item, status of the content item, activity of other users previously or currently viewing the content item. Such an overlay can warn a user that changes might be lost because another user is currently editing the content item.

[89] In some embodiments, one or more of the services or storages/databases discussed above can be accessed using public or private application programming interfaces.

[90] Certain software applications can access content storage 142 via an API on behalf of a user. For example, a software package such as an application running on client device 150, can programmatically make API calls directly to content management system 110 when a user provides authentication credentials, to read, write, create, delete, share, or otherwise manipulate content.

[91] A user can view or manipulate content stored in a user account via a web interface generated and served by web interface service 124. For example, the user can navigate in a web browser to a web address provided by content management system 110. Changes or updates to content in the content storage 142 made through the web interface, such as uploading a new version of a content item, can be propagated back to other client devices associated with the user's account. For example, multiple client devices, each with their own client software, can be associated with a single account and content items in the account can be synchronized between each of the multiple client devices.

[92] Client device 150 can connect to content management system 110 on behalf of a user. A user can directly interact with client device 150, for example when client device 150 is a desktop or laptop computer, phone, television, internet-of-things device, etc. Alternatively or additionally, client device 150 can act on behalf of the

user without the user having physical access to client device 150, for example when client device 150 is a server.

[93] Some features of client device 150 are enabled by an application installed on client device 150. In some embodiments, the application can include a content management system specific component. For example, the content management system specific component can be a stand-alone application 152, one or more application plug-ins, and/or a browser extension. However, the user can also interact with content management system 110 via a third-party application, such as a web browser, that resides on client device 150 and is configured to communicate with content management system 110. In various implementations, the client-side application 152 can present a user interface (UI) for a user to interact with content management system 110. For example, the user can interact with the content management system 110 via a file system explorer integrated with the file system or via a webpage displayed using a web browser application.

[94] In some embodiments, client application 152 can be configured to manage and synchronize content for more than one account of content management system 110. In such embodiments client application 152 can remain logged into multiple accounts and provide normal services for the multiple accounts. In some embodiments, each account can appear as folder in a file system, and all content items within that folder can be synchronized with content management system 110. In some embodiments, client application 152 can include a selector to choose one of the multiple accounts to be the primary account or default account.

[95] While content management system 110 is presented with specific components, it should be understood by one skilled in the art, that the architectural configuration of system 100 is simply one possible configuration and that other configurations with more or fewer components are possible. Further, a service can have more or less functionality, even including functionality described as being with another service. Moreover, features described herein with respect to an embodiment can be combined with features described with respect to another embodiment.

[96] While system 100 is presented with specific components, it should be understood by one skilled in the art, that the architectural configuration of system 100

is simply one possible configuration and that other configurations with more or fewer components are possible.

Client Synchronization Service

[97] FIG. 2 shows an example of a client synchronization service 156, in accordance with some embodiments. According to some embodiments, client synchronization service 156 may be implemented in the client device of FIG. 1. However, in other embodiments, client synchronization service 156 may be implemented on another computing device. Client synchronization service 156 is configured to synchronize changes to content items between a content management system and the client device on which client synchronization service 156 runs.

[98] Client synchronization service 156 may include file system interface 205, server interface 210, tree storage 220, planner 225, and scheduler 230. Additional or alternative components may also be included. High level descriptions of client synchronization service 156 and its components are discussed below with respect to FIG. 2. However, further details and embodiments of client synchronization service 156 and its components are discussed throughout.

[99] File system interface 205 is configured to process changes to content items on the local filesystem of the client device and update the local tree. For example, file system interface 205 can be in communication with client synchronization service 156 of FIG. 1 to detect changes to content items on the local filesystem of the client device. Changes may also be made and detected via client application 152 of FIG. 1. File system interface 205 may make updates to the local tree. The updates to the local tree may be made based on the changes (new, deleted, modified, copied, renamed, or moved content items) to content items on the client device.

[100] Server interface 210 is configured to aid in the processing of remote changes to content items at a remote storage of the content management system and updating of the remote tree. For example, server interface 210 can be in communication with server synchronization service 112 of FIG. 1 to synchronize changes to content items between client device 150 and content management system 110. Changes (new, deleted, modified, copied, renamed, or moved content items) to content items at content management system 110 may be detected and updates may be made to the remote tree to reflect the changes at content management system 110.

[101] Tree storage 220 is configured to store and maintain the tree data structures used by client synchronization service 156. For example, tree storage 220 may store the local tree, the sync tree, and the remote tree. According to some embodiments, tree storage 220 may store the tree data structures in persistent memory (e.g., a hard disk or other secondary storage device) as well as in main memory (e.g., RAM or other primary storage device) in order to reduce latency and response time. For example, on start-up of the client device or client synchronization service 156, the tree data structures may be retrieved from persistent memory and loaded into main memory. Tree storage 220 may access and update the tree data structures on main memory and, before the client device or client synchronization service 156 is shut down, tree storage 220 may store the updated tree data structures on persistent memory. Because main memory is expensive in cost and often limited in size on most client devices, additional technological improvements are implemented to decrease the footprint of the tree data structures on main memory. These technological solutions are described further below.

[102] Planner 225 is configured to detect differences between the server state associated with the content management system and the file system state associated with the client device based on the state of the tree data structures. For example, planner 225 may determine if there is a difference between the remote tree and the sync tree. A difference between the remote tree and the sync tree indicates that an action performed remotely on one or more content items stored at the content management system has caused the server state and the file system state to become out of sync. Similarly, planner 225 may also determine if there is a difference between the local tree and the sync tree. A difference between the local tree and the sync tree indicates that an action performed locally on one or more content items stored on the client device has caused the server state and the file system state to become out of sync. If a difference is detected, planner 225 generates a set of operations that synchronize the tree data structures.

[103] In some scenarios, a set of operations generated based on a difference between the remote tree and the sync tree and a set of operations generated based on a difference between the local tree and the sync tree may conflict. Planner 225 may also be configured to merge the two sets of operations into a single merged plan of operations.

[104] Scheduler 230 is configured to take the generated plan of operations and manage the execution of those operations. According to some embodiments, scheduler 230 converts each operation in the plan of operations into a series of one or more tasks that need to be executed in order to perform the operation. In some scenarios, some tasks may become out dated or no longer relevant. Scheduler 230 is configured to identify those tasks and cancel them.

Tree Data Structures

[105] FIG. 3 shows an example of tree data structures, in accordance with various embodiments. The tree data structures may be stored at the client device and managed by a client synchronization service such as client synchronization service 156 in FIG. 2. In FIG. 3, the tree data structures are shown including remote tree 310, sync tree 330, and local tree 350.

[106] Remote tree 310 represents a server state or the state of content items stored remotely from the client device (e.g., on a server of the content management system). Local tree 350 represents a file system state or the state of the corresponding content items stored locally on the client device. Sync tree 330 represents a merge base for the local tree and the remote tree. The merge base may be thought of as a common ancestor of the local tree and the remote tree or a last known synced state between the local tree and the remote tree.

[107] Each tree data structure (e.g., remote tree 310, sync tree 330, or local tree 350) may include one or more nodes. Each node may have one or more child nodes and the parent-child relationship is represented by an edge. For example, remote tree 310 includes nodes 312 and 314. Node 312 is a parent of node 314 and node 314 is a child of node 312. This parent-child relationship is represented by edge 316. A root node, such as root node 312, does not have a parent node. A leaf node, such as node 314, does not have a child node.

[108] Each node in a tree data structure may represent a content item (e.g., a file, document, folder, etc.). For example, root node 312 may represent the root folder associated with the content management system and node 314 may represent a file (e.g., a text file named "Foo.txt") located in that root folder. Each node in a tree data structure may contain data such as, for example, a directory file identifier ("DirFileID") specifying the file identifier of a parent node of the content item, a file

name for the content item, a file identifier for the content item, and metadata for the content item.

[109] As described above, a client synchronization service may determine that the server state and the file system state of the client device are synchronized when all 3 trees (e.g., remote tree 310, sync tree 330, and local tree 350) are identical. In other words, the trees are synchronized when their tree structures and the relationships that they express are identical and the data contained in their nodes are identical as well. Conversely, the trees are not synchronized if the 3 trees are not identical. In the example scenario illustrated in FIG. 3, remote tree 310, sync tree 330, and local tree 350 are shown as being identical and synchronized and, as a result, the server state and the file system state are synchronized.

Tracking Changes Using Tree Data Structures

[110] FIG. 4 shows an example of tree data structures, in accordance with various embodiments. As with the tree data structures shown in FIG. 3, the tree data structures shown in FIG. 4 (including remote tree 410, sync tree 430, and local tree 450) may be stored at the client device and managed by a client synchronization service such as client synchronization service 156 in FIG. 2. In FIG. 4, the tree data structures are shown.

[111] FIG. 4 shows a scenario after a previously synchronized state, such as the scenario illustrated in FIG. 3, additional actions are performed on the content items represented in the trees to modify the content items such that the trees are no longer in sync. Sync tree 430 maintains a representation of the previously known synchronized state and may be used by the client synchronization service to identify the differences between the server state and the file system state as well as generate operations for the content management system and/or the client device to perform to converge so that the server state and the file system state are synchronized.

[112] For example, a user (the same user as the user associated with the client device or a different user with access to the content item) may make modifications to the “foo.txt” content item stored by the content management system. This content item is represented by node 414 in remote tree 410. The modification shown in the remote tree 410 is a removal (e.g., a removal of the content item from a space managed by the content management system) or delete of the foo.txt content item. These

modifications may be performed, for example, on another client device and then synchronized to the content management system or performed through a web browser connected to the content management system.

[113] When the change is made on the content management system, the content management system generates modification data specifying the change made and transmits the modification data to the client synchronization service on the client device. For example, using a push model where the content management system may transmit or “push” changes to the client device unilaterally. In other implementations, a pull model where the server sends the changes in response to a request by the client device. Additionally, a hybrid model involving a long pull where the client device initiates the requests but keeps the connection open for a period of time so the content management system can push additional changes as needed while the connection is live. The client synchronization service updates the remote tree representing the server state for the content items stored by the content management system based on the modification data. For example, in remote tree 410, node 414 representing the foo.txt content item is shown as deleted.

[114] The client synchronization service may identify a difference between remote tree 410 and sync tree 430 and, as a result, determine that a modification of the content items at the content management system has caused the server state and the file system state to no longer be in sync. The client synchronization service may further generate and execute a set or sequence of operations for the content items stored on the client device that are configured to converge the server state and the file system state so that they will be in sync.

[115] Additionally or alternatively, a user (the same user as the user associated with modifications at the content management system or a different user with access to the content item) may make modifications to the content items stored locally on the client device that are associated with the content management system. For example, the user may add a folder “/bar” to the “/root” folder and add a “Hi.doc” document to the “/bar” folder.

[116] When the change is made on the client device, the client device (e.g., client synchronization service 156 or client application 152 of FIG. 1) generates modification data specifying the change made. The client synchronization service

updates the local tree representing the file system state for the content items stored on the client device based on the modification data. For example, in local tree 450, node 452 and node 454 are shown as added. Node 452 and node 454 represent the “/bar” folder and the “Hi.doc” document respectively.

[117] The client synchronization service may identify a difference between local tree 450 and sync tree 430 and, as a result, determine that a modification of the content items at the client device has caused the server state and the file system state to no longer be in sync. The client synchronization service may further generate a set or sequence of operations for the content items stored by the content management system that are configured to converge the server state and the file system state so that they will be in sync. These operations may be transmitted to the content management system for execution.

[118] As seen in FIG. 4, modifications to content items stored on the client device and content items stored by the content management system may occur at substantially the same time or within a particular time period. These modifications can be reflected in the tree data structures and used by the client synchronization service to generate operations for the client device and for the content management system in parallel. In other scenarios, however, modifications may not necessarily occur within the same time period and operations may be generated in an as-needed manner. Furthermore, although FIG. 4 illustrates scenarios for adding content items and deleting content items, other types of modifications such as, editing, renaming, copying, or moving content items are also supported.

[119] According to various embodiments, identifying a difference between two tree data structures and generating operations may involve checking each node in both tree data structures and determining whether an action has been performed on the node. The actions may include, for example, the addition of the node, the deletion of the node, the editing of the node, or the moving of the node. These actions may then be used to generate the operations configured to converge the server state and the file system state.

[120] For example, if the two tree data structures are a sync tree and a remote tree, the client synchronization service may identify each node in the sync tree by, for example, requesting the file identifiers of all nodes in the sync tree. For each node or

file identifier for the node in the sync tree, the client synchronization service may determine if the node or file identifier is also in the remote tree. A node or file identifier in the sync tree that is not found in the remote tree may indicate that the node has been deleted from the server state that is represented by the remote tree. Accordingly, the client synchronization service may determine that a delete action has occurred on the remote tree. If the node or file identifier for the node is found in the remote tree, the client synchronization service may check whether the node in the remote tree has been edited or moved.

[121] To determine whether the node in the remote tree has been edited with respect to the node in the sync tree, the client synchronization service may compare the metadata for the node in the sync tree with the metadata for the corresponding node (e.g., the node with the same file identifier) in the remote tree. The metadata may include information that may be used to determine whether the content item represented by the node has been edited. For example, the metadata may include one or more hash values that are generated based on the data in the content item or a portion thereof. The metadata may additionally or alternatively include a size value, a last modified value, or other value for the content item. The metadata for the node in the sync tree may be compared with the metadata for the node in the remote tree. If the metadata do not match, an edit of the content item may have been edited in the server state represented by the remote tree. Accordingly, the client synchronization service may determine that an edit action has occurred for the node on the remote tree. If the metadata matches, no edit may have occurred.

[122] To determine whether the node in the remote tree has been moved, the client synchronization service may compare the location for the node in the sync tree with the location for the corresponding node (e.g., the node with the same file identifier) in the remote tree. The location may include, for example, a path where the node is located, a file name, and/or a directory file identifier (“DirFileID”) specifying the file identifier of the node’s parent. If the locations match, no move may have occurred. On the other hand, if the locations do not match, a move of the content item may have occurred in the server state represented by the remote tree. Accordingly, the client synchronization service may determine that a move action has occurred for the node on the remote tree.

[123] To determine whether a node has been added to the remote tree, the client synchronization service may identify any nodes or file identifiers in the remote tree that are not found in the sync tree. If a node or file identifier is found in the remote tree and not found in the sync tree, the client synchronization service may determine that an add action of this node has occurred on the remote tree representing the server state.

[124] Although the example above is described with respect to the sync tree and the remote tree, in other embodiments, a similar process may occur with the sync tree and a local tree in order to identify a difference between the sync tree and the local tree and determine which actions have occurred on the local tree representing the file system state.

Synchronization Using Tree Data Structures

[125] FIG. 5 shows an example method for synchronizing a server state and a file system state using tree data structures, in accordance with various embodiments of the subject technology. Although the methods and processes described herein may be shown with certain steps and operations in a particular order, additional, fewer, or alternative steps and operations performed in similar or alternative orders, or in parallel, are within the scope of various embodiments unless otherwise stated. The method 500 may be implemented by a system such as, for example, client synchronization service 156 of FIG. 2, running on a client device.

[126] The system is configured to identify a difference between a remote tree representing a server state for content items stored by the content management system, a local tree representing the file system state for the corresponding content items stored on the client device, and a sync tree representing a known synced state between the server state and the file system state. Based on these differences, a set of operations may be generated that, if executed, are configured to converge the server state and the file system state towards a synchronized state where the three tree data structures would be identical.

[127] For example, at operation 505, the system may receive modification data for content items stored by a content management system or on a client device. The modification data may be used to update a remote tree or a local tree at operation 510.

[128] The modification data specifies what changes occurred to one or more content items associated with a content management service. Accordingly, the modification data may be received from the content management system or from the client device (e.g., from client application 152 running on client device 150 in FIG. 1). Modification data received from the content management system may be referred to as server modification data. Server modification data specifies what changes are done to one or more content items by the content management system and may be used to update the remote tree at operation 510. Modification data received from the client device may be referred to as client modification data. Client modification data specifies what changes are done to one or more content items on the client device and may be used to update the local tree at operation 510.

[129] At operation 515, the system may determine whether a server state for content items stored by the content management system and a file system state for the content items stored on the client device are in sync. Because the local tree and the remote tree are representative of the file system state and the server state and are continually being updated to track changes that occur at the content management system and the client device, determining whether the server state and the file system state are in sync may be done by comparing the local tree and/or the remote tree to the sync tree to find differences between the trees. This process of finding differences between the trees is sometimes referred to as “diffing” the trees.

[130] According to some embodiments and scenarios, determining whether the server state and the file system state are in sync may include one or more of identifying differences between the remote tree and the sync tree and/or identifying differences between the local tree and the sync tree. Differences between the remote tree and sync tree may indicate the occurrence of changes to content items stored by the content management system that may not be reflected at the client device. Similarly, differences between the local tree and sync tree may indicate the occurrence of changes to content items stored at the client device that may not be reflected at the content management system.

[131] If there are no differences between the trees, the server state and the file system state are in sync and no synchronization actions are needed. Accordingly, the method may return to operation 505 and await new modification data. On the other

hand, if differences are detected, the system may generate a set of operations configured to converge the server state and the file system state at operation 520.

[132] The set of operations generated depends on the one or more differences that are detected. For example, if the difference between two trees is an added content item, the generated set of operations may include retrieving the added content item and adding it. If the difference between two trees is a deletion of a content item, the generated set of operations may include deleting the content item. According to some embodiments, the set of operations may also include a number of checks to ensure tree constraints are maintained. As will be described further below, the set of operations may conflict with the current state of the server state, the file system state, or other operations that are pending execution. Accordingly, the system may also resolve these conflicts before proceeding.

[133] As noted above, if there are differences between the remote tree and sync tree, changes to content items stored by the content management system may have occurred that may not be reflected at the client device. Accordingly, in this scenario, the system may generate a client set of operations configured to operate on the content items stored on the client device to converge the server state and the file system state and this client set of operations may be provided to the client device for execution at operation 525.

[134] Similarly, if there are differences between the local tree and sync tree, changes to content items stored at the client device may have occurred that may not be reflected at the content management system. Accordingly, in this scenario, the system may generate a server set of operations configured to operate on the content items stored by the content management system to converge the server state and the file system state and this server set of operations may be provided to the content management system for execution at operation 525. In some cases, both cases may be true and a client set of operations and a server set of operations may be generated and provided to their intended recipients at operation 525.

[135] Once the set(s) of operations are provided to the intended recipient(s), the method may return to operation 505 and await new modification data. The set(s) of operations may provide one or more steps towards the convergence of the server state and the file system state or provide all steps needed to sync the server state and the

file system state. For example, the content management system may receive the server set of operations and execute the server set of operations on content items stored by the content management system. This execution of the server set of operations causes changes to the content items stored by the content management system, which are detected and specified in server modification data, which is transmitted back to the system. The system may then update the remote tree and determine whether the server state and the file system state are in sync.

[136] The client device may receive the client set of operations and execute the client set of operations on content items stored on the client device. This execution of the client set of operations causes changes to the content items stored on the client device, which are detected and specified in client modification data, which is passed to the system. The system may then update the local tree and determine whether the server state and the file system state are in sync. These operations of method 500 may continue until the server state and the file system state are in sync.

[137] The operations of method 500 are described with respect to a client side and a server side (e.g., a local tree and a remote tree, a file system state and a server state, a client set of operations and a server set of operations, client modification data and server modification data). In various embodiments the operations associated with the two sides may occur in parallel, in sequence, in isolation of the other side, or a combination.

[138] As will be discussed in further detail, in accordance with some embodiments, before the operations are provided for execution, the system may check the operations to determine whether they comply with a set of rules or invariants. If an operation violates a rule, the system executes a resolution process associated with the violation of the rule.

[139] Additionally, in accordance with some embodiments, the system (e.g., scheduler 230 of client synchronization service 156 in FIG. 2) may manage the execution of the set of operations. For example, each operation in the set of operations may be associated with a task, an execution thread, series of steps, or instructions. The system may be configured to execute the task, thread, step, or instructions and interface with the client device and/or the content management

system to execute the set of operations and converge the server state and the file system state.

Conflict Handling

[140] As described above with respect to FIG. 5, differences between a sync tree and a remote tree are identified and used to generate a client set of operations configured to converge the server state and the file system state. However, in some cases, the client set of operations may conflict with the current state of a local tree. Similarly, differences between the sync tree and the local tree are identified and used to generate a server set of operations configured to converge the server state and the file system state. However, the server set of operations may conflict with the current state of the remote tree. Additionally or alternatively, the client set of operations and the server set of operations may conflict with one another or violate another rule or invariant maintained by the system. Accordingly, various embodiments of the subject technology provide additional technical improvements by resolving these conflicts.

[141] For example, planner 225 in client synchronization service 156 of FIG. 2 may identify an operation in a set of operations (e.g., the client set of operations or the server set of operations) that conflicts with a rule. Each rule used to identify a conflict may also be associated with a resolution for the conflict. The client synchronization service may update the set of operations based on the resolution for the conflict or resolve the conflict by performing operations associated with the resolutions for the conflict before providing the set of operations for execution.

[142] FIG. 6 shows an example method 600 for resolving conflicts when synchronizing a server state and a file system state using tree data structures, in accordance with various embodiments of the subject technology. Although the methods and processes described herein may be shown with certain steps and operations in a particular order, additional, fewer, or alternative steps and operations performed in similar or alternative orders, or in parallel, are within the scope of various embodiments unless otherwise stated. The method 600 may be implemented by a system such as, for example, client synchronization service 156 of FIG. 2, running on a client device.

[143] The system may receive a set of operations configured to converge a server state and a file system state at operation 620. The set of operations may be, for

example, the client set of operations, the server set of operations, or a combined set of operations generated and described with respect to the method 500 of FIG. 5.

[144] At operation 650, the system identifies one or more violations in the set of operations based on a set of rules. The set of rules may be stored by client synchronization service 156 in FIG. 2 and specify a number of constraints, invariants, or conflicts for operations that are to be resolved. The set of rules may be applied to the tree data structures and help control synchronization behavior. Each rule in the set of rules may also be associated or otherwise linked to a resolution to a violation of that rule. For example, the resolution may include an alteration of one or more operations in the set of operations, a removal off one or more operations, an addition of one or more operations, one or more additional actions to the server state or the file system state, or a combination of actions.

[145] For each operation in a set of operations, the system may determine whether any rule in the set of rules is violated. If a rule is violated, the system identifies a resolution of the violation and, at operation 655, performs the resolution. The resolution may include actions such as modifying one or more operations in the set of operations, a removing or adding one or more operations, or additional actions on the server state or the file state.

[146] Once the resolution actions are performed, the system may generate a resolved or rebased set of operations based on the resolution and the set of operations at operation 660 and, at operation 665, provide the resolved set of operations to the appropriate entity for execution. For example, the resolved set of operations may be provided to scheduler 230 of client synchronization service 156 in FIG. 2 for managed execution. Alternatively, if the set of operations is a client set of operations, the resolved set of operations may be provided to the client device. If the set of operations is a server set of operations, the resolved set of operations may be provided to the content management service. Additionally, the method 600 of FIG. 6 may be performed on client set of operations and server set of operations in sequence, in parallel, or in various different orders.

[147] According to some embodiments, each type of operation may be associated with the same or a different set of rules. For example, operation types may include, for example, adding a content item, deleting a content item, editing a content item,

moving a content item, renaming a content item, etc. The set of operations may consist of operations each belonging to one of the operation types above. Each operation type may be associated with a specific set of rules.

[148] For illustrative purposes, a set of rules for an “Add” operation type may include rules such as file identifiers for content items must be unique in a tree (e.g., no two nodes in a tree may have the same file identifier), a directory file identifier (“DirFileID”) specifying the file identifier of a parent node of the content item must exist in the opposite tree data structure, and a DirFileID and file name combination for a content item are not used in the opposite tree.

[149] Opposite tree, as used here, refers to the tree data structure that represents the state of the opposing entity. For example, a client set of operations configured to operate on the client device and the resulting changes to the file system on the client device will be reflected in the local tree. Accordingly, the opposite tree for the client set of operations is the remote tree. Similarly, a server set of operations is configured to be transmitted to the content management system to be executed and the resulting changes to the server state will be reflected in the remote tree. Accordingly, the opposite tree for the server set of operations is the local tree.

[150] FIG. 7 shows an example of tree data structures illustrating a violation of a rule for an add operation, in accordance with various embodiments. The tree data structures include remote tree 710, sync tree 750, and local tree 770. When referencing the local tree 770, the remote tree 710 may be considered the opposite tree. On the other hand, when referencing the remote tree 710, the local tree 770 may be considered the opposite tree. FIG. 7 illustrates a set of operations adding the content item represented by node 712 in remote tree 710. For example, a client synchronization service may compare remote tree 710 with sync tree 750, identify the differences, and generate a set of operations that includes the addition of node 712. Node 712 is associated with a FileID of 4, a DirFileID of 3 (which references parent node 714, which is node 712’s parent), and a file name of “Hi.” Parent node 714 is associated with a FileID of 3, a DirFileID of 1 (which references root node 716, which is node 714’s parent), and a file name of “Foo.”

[151] The client synchronization service may perform the method 600 of FIG. 6 and determine that the add operation for node 712 violates the “a directory file identifier

(“DirFileID”) of the content item must exist in the opposite tree data structure” rule for “add” operation types. This is illustrated in FIG. 7 by the local tree 770 not having a node with a file ID of 3, which references parent node 714 of node 712. This may occur when, for example, after differences between remote tree 710 and sync tree 750 are determined and a set of operations is generated, the “Foo” node corresponding to node 714 is removed from the opposite tree.

[152] The resolution associated with this rule may include deleting the node missing from local tree 770 from sync tree 750 to synchronize sync tree 750 and local tree 770 and rediffing (e.g., finding the difference between) remote tree 710 and sync tree 750. In the scenario illustrated in FIG. 7, node 754 in sync tree 750 would be removed 758 and diffing operations would commence to identify differences between remote tree 710 and sync tree 750. This would result in the inclusion of an add operation of node 714 as well as an add operation for node 712 in the set of operations.

[153] Similarly, a violation of the “file identifiers for content items must be unique in a tree” rule for “add” operation types may be resolved by operations including requesting, from the content management system, a new file ID for the node being added and using the new file ID when adding the node. A violation of the “DirFileID and file name combination for a content item are not used in the opposite tree” rule for “add” operation types may be resolved by operations including checking via the metadata associated with the two nodes whether the content items are the same. If the content items are the same, it is likely that the content item being added has already been added in other actions. If the content items are not the same, the file name for the content item being added can be renamed. For example, the file name for the content item being added can be appended with the text “(conflicted version).”

Incremental Planner

[154] Although the various tree data structures shown in FIGS. 3, 4, and 7 contain a relatively small number of nodes and are relatively simple in structure, the tree data structures supported by the system may be much larger and complex with multiple levels and potentially large number of nodes at each level. Accordingly the memory usage required to store the tree data structures during operation may be quite large and the computing time and resources required to operate on the tree data structures may be quite large. For example, finding differences between a remote tree and a sync tree

and/or a local tree and the sync tree and generating operations needed to converge the remote tree and the sync tree and/or the local tree and the sync tree may require a large amount of memory, time, and other computing resources.

[155] Unfortunately, these computing resources are limited. For example, a client device may have a limited amount of available memory and the length of time needed to diff trees and generate operations may hinder the usability of the client device, the client application, or the content management services provided by the content management system. Furthermore, the more time needed to converge the server state and the file system state, the more likely that intervening changes to either state may render the set of operations being computed or executed and/or the target sync state out of date. Accordingly, various embodiments of the subject technology provide additional technical improvements by incrementally converging the server state and the file system state along with the tree data structures that represent them.

[156] FIG. 8 shows an example method 800 for incrementally converging a server state and a file system state, in accordance with various embodiments of the subject technology. Although the methods and processes described herein may be shown with certain steps and operations in a particular order, additional, fewer, or alternative steps and operations performed in similar or alternative orders, or in parallel, are within the scope of various embodiments unless otherwise stated. The method 800 may be implemented by a system such as, for example, client synchronization service 156 of FIG. 2, running on a client device.

[157] At operation 805, the system may receive modification data that may be used to update either a remote tree or a local tree. For example, server modification data may be received from a content management system that specifies modifications or other actions (e.g., an edit, add, delete, move, or rename) associated with one or more content items stored by the content management system. The server modification data may be used to update the remote tree, which represents the server state of content items stored by the content management system. Similarly, client modification data may be received from the client device (e.g., a client application) and specify modifications or other actions associated with one or more content items stored on the client device. The client modification data may be used to update the local tree, which represents the file system state of content items stored on the client device.

[158] Based on the received modification data specifying modifications associated with content items, the system may identify nodes that correspond to the modified content items and add the nodes to a list of modified content items (e.g., add the file identifier associated with the nodes to the list of modified content items) at operation 810. Operations 805 and 810 may continuously occur for some time before the system proceeds to the next stage of the method 800. For example additional modification data may be received and used to update the trees managed by the system and add nodes to the list of modified content items.

[159] In order to incrementally converge the server state and the file system state, the system takes each node in the list of modified content items and determines how the node was modified (e.g., which actions are associated with the node) at operation 815. In some embodiments, the modification data may specify the modification to the node. However, in other embodiments, the system may determine the modifications to the node based on a comparison of the remote tree with the sync tree and/or a comparison of the local tree with the sync tree. For example, the modifications may include the addition of the node, the deletion of the node, the editing of the node, or the moving of the node.

[160] For each node or file identifier for the node in the list of modified content items, the system may perform a series of checks to determine what, if any, modifications were performed on the node. For example, the system may determine whether the file identifier is in the sync tree but not in the remote tree. A file identifier in the sync tree that is not found in the remote tree may indicate that the node has been deleted from the server state that is represented by the remote tree. Accordingly, the client synchronization service may determine that a delete modification on the node has occurred on the remote tree. Similarly, the system may also determine whether the file identifier is in the sync tree but not in the local tree. A file identifier in the sync tree that is not found in the local tree may indicate that the node has been deleted from the file system state that is represented by the local tree. Accordingly, the client synchronization service may determine that a delete modification on the node has occurred on the local tree.

[161] To determine whether an edit modification has been performed on the node, the system may compare the metadata for the node in the sync tree with the metadata for the corresponding node (e.g., the node with the same file identifier) in the remote

tree and/or the local tree. The metadata may include information that may be used to determine whether the content item represented by the node has been edited. For example, the metadata may include one or more hash values that are generated based on the data in the content item or a portion thereof. The metadata may additionally or alternatively include a size value, a last modified value, or other value for the content item. If the metadata do not match, an edit of the content item may have been edited in the server state represented by the remote tree and/or the file system state represented by the local tree. Accordingly, the system may determine that an edit action has occurred for the node on the remote tree and/or the local tree.

[162] To determine whether the node in the remote tree has been moved, the system may compare the location for the node in the sync tree with the location for the corresponding node (e.g., the node with the same file identifier) in the remote tree and/or the local tree. The location may include, for example, a path where the node is located, a file name, and/or a directory file identifier (“DirFileID”) specifying the file identifier of the node’s parent. If the locations match, no move may have occurred. On the other hand, if the locations do not match, a move of the content item may have occurred in the remote tree or the local tree. Accordingly, the client synchronization service may determine that a move action has occurred for the node on the remote tree and/or the local tree.

[163] To determine whether a node has been added to the remote tree, the system may determine if the file identifier in the list of modified content items is in the remote tree or in the local tree, but not in the sync tree. If the file identifier is found in the remote tree or the local tree and not found in the sync tree, the system may determine that an add modification for this node has occurred.

[164] Once the one or more modifications to the nodes in the list of modified content items are determined, the system may determine whether any of those modifications have dependencies at operation 820. As will be illustrated further with respect to FIG. 9, a modification on a node has a dependency when, for example, the modification cannot execute without another modification occurring first.

[165] If the modification does not have a dependency, the system adds the modification to an unblocked list of actions at operation 825. If the modification has a dependency, the modification is blocked for the time being at operation 830 and

cannot be executed without another modification being processed first. Accordingly the process returns to operation 805 to await further modifications. After each of the modifications are processed, the system may clear the file identifiers associated with the modifications from the list of modified content items.

[166] FIG. 9 shows an example of tree data structures, in accordance with various embodiments. The tree data structures shown in FIG. 9 may be stored at the client device and managed by a system such as client synchronization service 156 in FIG. 2. For the purpose of illustration, only remote tree 910 and sync tree 950 are shown in FIG. 9 and described. Similar operations and description may also be applied to a local tree as well.

[167] Remote tree 910 includes root node 912 with a file identifier of 1, node 914 with a file identifier of 5 and file name of “Foo,” node 916 with a file identifier of 6 and file name of “Bar,” and node 918 with a file identifier of 7 and file name of “Bye.” Sync tree includes root node 952 with a file identifier of 1.

[168] Based on the tree data structures shown in FIG. 9, the system may have identified that nodes with file identifiers of 5, 6, and 7 have been modified at operation 810 and added the nodes to the list of modified content items, as illustrated by reference 980 in FIG. 9. At operation 815, the system determines the list of modifications to nodes in the list of modified content items. As is seen by the comparison of remote tree 910 and sync tree 950, nodes 914, 916, and 918 have been added to remote tree 910. More specifically, as illustrated by reference 982 in FIG. 9, node 916 with file identifier 6 and name “Bar” has been added as a child to node 914 with file identifier 5. This is represented by the “Add(6, 5, Bar)” entry in reference 982. Node 918 with file identifier 7 and name “Bye” has been added as a child to node 914 with file identifier 5. This is represented by the “Add(7, 5, Bye)” entry in reference 982. Node 914 with file identifier 5 and name “Foo” has been added as a child to root node 912 with file identifier 1. This is represented by the “Add(5, /root, Foo)” entry in reference 982.

[169] At operation 820, the system determines that the add modification of node 914 does not have a dependency and, as a result, is unblocked. Accordingly, the system adds the modification associated with node 914 (e.g., the modification represented by the “Add(5, /root, Foo)” entry in reference 982) to an unblocked list of actions at

operation 825. This is seen in references 984 in FIG. 9. On the other hand, the modifications for nodes 916 and 918 represented by the “Add(6, 5, Bar)” and the “Add(7, 5, Bye)” entries in reference 982 are dependent on the modification represented by the “Add(5, /root, Foo)” occurring first. In other words, node 916 and/or node 918 cannot be added until node 914 is added. Accordingly, these modifications are included in a blocked list of actions illustrated by reference 986 in FIG. 9.

[170] Returning to the method 800 of FIG. 8, at operation 835, the system may select a set of modifications from the unblocked list of actions and generate a set of operations based on the selected set of modifications. The set of operations is configured to converge the server state and the file system state. The set of operations generated depends on the selected set of modifications from the unblocked list. For example, if the selected set of modifications includes the add modification associated with node 914 (e.g., the modification represented by the “Add(5, /root, Foo)” entry in reference 984) in FIG. 9, the generated set of operations may include retrieving the added content item from the content management system and adding it to the local file system of the client device.

[171] According to some embodiments, the system may select all modifications from the unblocked list of actions to generate one or more sets of operations. However, in some scenarios, the number of modifications in the unblocked list may be quite high and the computing resources (e.g., memory and processing time) needed to process all of the modifications is substantial. In order to reduce these technological burdens, the system may select a smaller set of the modifications in the unblocked list of actions in order to process incrementally. For example, the system may select the first or top X number or percent of modifications to generate operations. In further iterations of the process, the remaining modifications in the unblocked lists may be processed.

[172] In some embodiments, the modifications in the unblocked list may be ranked for processing. The modifications may be ranked based on, for example, a modification type (e.g., delete modifications are prioritized over add modifications), metadata associated with the modification (e.g., add modifications of content items of smaller size are prioritized over add modifications of content items of larger size,

delete modifications of content items of larger size are prioritized over delete modifications of content items of smaller size, etc.).

[173] These rank rules may be stored by the system and may be designed to achieve various performance goals for content synchronization. For example, delete modifications may be prioritized over add modifications in order to free as much of potentially limited storage space for a user before new content items may be added. Adding of smaller content items may be prioritized over larger content items in order to provide as much progress with respect to the number of content items added as soon as possible.

[174] At operation 835, the system may provide the set of operations to the content management system and/or the client device. As noted above, modifications associated with actions performed by the content management system may not be reflected at the client device. Accordingly, in this scenario, the system may generate a client set of operations configured to operate on the content items stored on the client device to converge the server state and the file system state and this client set of operations may be provided to the client device for execution at operation 835.

[175] On the other hand, modifications associated with actions performed by the client device may not be reflected at the content management system. Accordingly, in this scenario, the system may generate a server set of operations configured to operate on the content items stored by the content management system to converge the server state and the file system state and this server set of operations may be provided to the content management system for execution at operation 835.

[176] In some cases, both cases may be true and a client set of operations and a server set of operations may be generated and provided to their intended recipients at operation 835. The set of operations may also include a number of checks to ensure tree constraints are maintained. For example, the set of operations may resolve various conflicts or constraints as discussed with respect to FIG. 6.

[177] Once the set(s) of operations are provided to the intended recipient(s), the method may return to operation 805 and await new modification data. For example, with respect to the scenario illustrated in FIG. 9, the set of operations may include retrieving the content item associated with node 914 from the content management system and adding it to the local file system of the client device. This would result in

the addition of a node corresponding to node 914 in the local tree (not shown in FIG. 9) and sync tree 950. On the next iteration of process 800 of FIG. 8, the add modifications of node 916 and node 918 represented by the “Add(6, 5, Bar)” and the “Add(7, 5, Bye)” entries in reference 982 are no longer blocked because their parent, node 914, has already been added to the sync tree. Accordingly, the add modifications of node 916 and node 918 represented by the “Add(6, 5, Bar)” and the “Add(7, 5, Bye)” entries in reference 982 may be added to the unblocked list of actions and used to generate one or more sets of operations configured to converge the server state and the file system state.

[178] The set(s) of operations may provide one or more steps for the incremental convergence of the server state and the file system state. Although implementing an incremental process may be more complex at times, the incremental process may achieve a reduction in processing time and reduction in the memory required. These and other initial technological improvements naturally lead to additional technological improvements. For example, because processing time is reduced, the likelihood of additional changes from the client device or the content management system making certain modifications obsolete or out of data is reduced as well.

[179] With respect to FIG. 9, various groupings of content items, modifications, actions, or file identifiers are described as lists for the purpose of illustration. Other types of data structures are also compatible. For example, the unblocked list of actions may be implemented as a B-tree data structure in order to keep data sorted and allow searches, sequential access, insertions, and deletions in logarithmic time.

Scheduler

[180] In some embodiments, a client synchronization service may generate a set or sequence of operations configured to converge the server state and the file system state and provide the operations to the content management system or client device for execution. However, in some scenarios, changes on the file system of the client device or on the content management system may cause the generated set of operations to become out of date or obsolete while the set of operations is in the process of executing. Various embodiments are directed to providing a technical solution to these and other technical problems. For example, the client synchronization service may be configured to monitor changes on the file system of

the client device or on the content management system and update the client device and/or content management system as needed. Furthermore, the client synchronization service may be configured to improve performance and reduce processing times by allowing for concurrent execution of operations.

[181] According to some embodiments, planner 225 of client synchronization service 156 shown in FIG. 2 may generate a plan or plan of operations that consists of an unordered set of operations. All operations within a plan have no dependencies and, as a result, are able to be executed concurrently in separate threads or in any order. The operations in the plan, according to some embodiments, are abstract instructions that may be taken by the content management system and/or the client device in order to converge the states and tree data structures. Example instructions may include a remote or local add of a content item, a remote or local delete of a content item, a remote or local edit of a content item, or a remote or local move of a content item.

[182] Scheduler 230 of client synchronization service 156 shown in FIG. 2 may be configured to receive the plan of operations from planner 225, manage the execution of the operations in the plan, determine if the plan has been updated or changed, and manage the execution of the updated or changed plan. For example, scheduler 230 may coordinate with file system interface 205 and server interface 210 to execute the tasks and steps needed to implement operations in the plan. This may include receiving confirmations from the file system or content management system or error handling activities such as handling retries when there is no network connectivity or when a content item is locked by some other application.

[183] Each operation may be implemented by a script or thread referred to as a task. The task coordinates the application of an associated operation and may include one or more steps needed to implement the operation. For example, a “local add operation” may indicate that a content item has been added to the local file system of the client device and, as a result, the content item should be added at the content management system in order to synchronize the server state and the file system state. Accordingly, the local add operation may be associated with a “local add task” that includes one or more steps needed to implement the local add operation. The steps may include one or more of notifying the content management system of the new content item, uploading the content item to the content management system in one or

more blocks of data, confirming that all blocks of data have been received by the content management system, making sure the content item is not corrupted, uploading metadata for the content item to the content management system, and committing the adding of the content item to the appropriate location at the content management system.

[184] A task may begin execution, suspend at well-defined points while waiting on the completion of other events, resume when the events have occurred, and eventually terminate. According to some embodiments, scheduler 230 is configured to cancel, regenerate, or replace tasks. For example, based on changes to the server state or the file system state, a task may become stale before it is executed and scheduler 230 may cancel the stale task before it is executed.

[185] As described above, planner 225 may generate a plan of operations based on a set of tree data structures (e.g., a remote tree, a sync tree, and a local tree). Over time, planner 225 continues to generate plans of operations based on the status of the tree data structures. If the tree data structures change to reflect the state of the server state and the file system state, planner 225 may also generate a new updated plan that differs from a previous plan. Scheduler 230 executes each plan of operations generated by the planner 225.

[186] In some scenarios, changes in the operations of a subsequent plan may cause unintended synchronization behaviors conflicts with an operation in the previous plan that is in the process of execution. For example, as operations in a first plan are being executed, one or more of the operations are canceled (or are not present) in the second plan. To illustrate, FIG. 10 shows an example scenario in which, at time t1, the server state represented by the remote tree and the file system state represented by the local tree are synchronized as shown by the remote tree, the sync tree, and the local tree all matching. Based on this synchronized state, planner 225 may generate a plan with no operations (e.g., an empty plan) at t1 or not generate a plan of operations.

[187] A user on the client device may delete content item A from the local file system or move content item A out of a folder managed by client synchronization service 156, which is reflected by the removal of node A from the local tree at time t2. Planner 225 may generate a plan that includes operation LocalDelete(A) based on the state of the tree data structures at time t2. Scheduler 230 may initiate the task or steps

required to implement the LocalDelete(A) operation. These steps may include transmitting instructions to the content management system to delete content item A.

[188] After instructions to delete content item A are transmitted to the content management system, the user on the client device may undo the delete of content item A or move content item A back to the previous location. The local tree is updated based on this new action at time t3 and planner may generate a new plan that is empty with no operations. Once again, the tree data structures match and the system is in a synchronized state at time t3.

[189] However, because instructions to delete content item A were transmitted to the content management system, the content management system deletes content item A from the server state. Although scheduler 230 may attempt to cancel the deletion of content item A, the instructions may have already been transmitted and completed by the content management system. This change in the server is communicated to client synchronization server 156, which updates the remote tree by deleting node A at time t4. Planner 225 could notice the change in the remote tree and the difference between the remote tree and the sync tree and determine that content item A was removed at the server state. Accordingly, planner 225 would create a plan with a RemoteDelete(A) operation at time t4. In an effort to synchronize the server state and the file system state, content item A would eventually be deleted from the client device and the local tree.

[190] Problematically, the removal of content item A from the server state, the generation of the RemoteDelete(A) operation, and the eventual removal of content item A from the file system state are all not intended and may cause further problems down the line for the user. Furthermore, in some cases, applications or processes may also access content items and unintentional synchronization behavior may cause a cascade of additional technical issues. Various embodiments are directed to preventing unintended consequences in synchronization of content items between a server state and a file system state.

[191] According to some embodiments, when canceling a task for a stale operation that is no longer in a plan of operations, scheduler 230 may wait for the cancelation to be completed before proceeding to initiate the execution of other tasks. For example, scheduler 230 may wait to receive confirmation of the cancelation from the client

device or the content management system before proceeding with other tasks. Scheduler 230 may determine whether the task has been initiated and if the task has not been initiated, scheduler may cancel the task and confirm that the task is no longer awaiting execution. If the task has been initiated, the confirmation may come from the client device or the content management system and notify the scheduler that all of the steps associated with the canceled task have been undone. According to some implementations, scheduler 230 does not allow for cancelation of a task once it has been initiated. This may be the case for all tasks or a certain subset of tasks or task types (e.g., a commit task that sends an update on the file system state to the content management system for synchronization with the server state).

[192] In order to improve performance and allow for concurrent execution of tasks as well as the cancelation of tasks, scheduler 230 may also be configured to manage the execution and cancelation of tasks based on differences between a first plan of operations and an updated second plan of operations. FIG. 11 shows an example Venn diagram 1100 representation of two plans of operations, in accordance with various embodiments of the subject technology. Planner 225 may generate a plan 1 1110 with a first set of operations, receive an update to the tree data structures, and generate an updated plan 2 1120 with a second set of operations.

[193] Plan 1 1110 and plan 2 1120 may share a number of common operations, which is represented by portion 1130 of the Venn diagram 1100. Plan 1 1110 and plan 2 1120 may also share a number of operations that are not in common. For example, operations in plan 1 1110 that are not in plan 2 1120 are stale and no longer current based on the update to the tree structures detected by planner 225. These stale operations of plan 1 1110 are represented by portion 1140 of Venn diagram 1100. New operations in plan 2 1120 that are not in plan 1 1110 are represented by portion 1150. Each of portions 1130, 1140, and 1150 which represent the differences and commonalities between plan 1 1110 and plan 2 1120 may include no operations or many operations depending on the updates to the server state and the file system state that are reflected in the tree data structures.

[194] Because the operations in portion 1140 are no longer in the most recent plan, scheduler 230 may cancel tasks associated with these operations. In order to prevent unintended synchronization behavior, tasks associated with operations in plan 2 that are not in plan 1 (e.g., in portion 1150) are postponed until the cancelation of tasks

associated with operation in portion 1140 is completed. However, because operations in each plan are configured to be able to be executed concurrently, tasks associated with operations in the intersection of plan 1 and plan 2 represented by portion 1130 may be executed concurrently with the cancelation of tasks associated with operation in portion 1140 without having to wait for their completion. By allowing for the concurrent cancelation of tasks associated with portion 1140 and the execution of tasks associated with portion 1130, more efficient use of available computing resources may be achieved as well as a reduction in processing time.

[195] FIG. 12 shows an example method for managing changes in plans of operations, in accordance with various embodiments of the subject technology. Although the methods and processes described herein may be shown with certain steps and operations in a particular order, additional, fewer, or alternative steps and operations performed in similar or alternative orders, or in parallel, are within the scope of various embodiments unless otherwise stated. The method 1200 may be implemented by a system such as, for example, client synchronization service 156 of FIG. 2, running on a client device.

[196] The system may be configured to receive updates from a content management system and/or the client device with regards to content items associated with a content management service. For example the system may receive server modification data for content items stored by a content management service and update, based on the server modification data, a remote tree. The remote tree represents the server state for content items stored by the content management system. The system may also receive client modification data for content items stored on the client device and update, based on the client modification data, a local tree. The local tree represents the file system state for content items stored on the client device.

[197] At operation 1205, the system may receive a first set of operations configured to converge a server state associated with the content management system and a file system state associated with the client device. For example, the system may identify differences between a sync tree and a remote tree or the sync tree and a local tree and generate the first set of operations based on any differences between the trees. The sync tree represents a known synced state between the server state and the file system state.

[198] The system may begin to implement the first set of operations. For example, in some cases, the operations are in a format ready to be transmitted to the content management system and/or the client device for execution. In other cases, the operations may be translated into one or more tasks, scripts, or execution threads that may be managed by the system. The system may interface with the content management system and/or the client device according to the tasks, scripts, or execution threads in order to converge the server state and the file system state.

[199] During this time, the system may continue to receive modification data from a content management system and/or the client device with regards to content items associated with the content management service. Based on the modification data, the system may update the remote tree or local tree and generate a second set of operations based on the updates to the tree data structures. At operation 1210, the system may receive the second set of operations.

[200] At operation 1215, the system identifies a first operation in the first set of operations that is not in the second set of operations, if any. If the system finds an operation in the first set of operations that is not in the second set of operations, this operation may be stale and out of date as a result of changes specified in the modification data. Accordingly, the system will initiate the cancelation of the first operation at operation 1220. The cancelation of the first operation may include a number of steps, a number of confirmation receipts for the steps, and a non-trivial amount of processing time.

[201] At operation 1225, the system identifies a second operation that is included in both the first set of operations and the second set of operations, if any. If the system finds an operation in both the first set of operations and the second set of operations, this operation may be still be valid notwithstanding changes specified in the modification data. Furthermore, since the operations in both sets of operations are configured to be able to be executed concurrently or in any order with respect to other operations in the set, the second operation can continue execution while the first operation is canceled. Accordingly, the system will initiate the execution of the second operation at operation 1230 without waiting for the first operation to complete cancelation.

[202] At operation 1235, the system identifies a third operation that is in the second set of operations, but not in the first set of operations, if any. If the system finds an operation in the second set of operations that is not in the first set of operations, this operation may be a new operation as a result of changes specified in the modification data. In order to prevent unintended consequences, the system will initiate a wait for the completion of the cancelation of the first operation. At operation 1240, the system may determine that the first operation has completed cancelation and, as a result, initiate the execution of the third operation at operation 1245.

Tree Data Structure Storage – Reduction of Filename Storage Space

[203] Client Synchronization Service 156 may store the tree data structures (e.g., the remote tree, the sync tree, and the local tree) on a persistent storage device such as, for example, a hard disk, solid state memory, or other types of computer readable media. In order to improve performance, reduce processing time, and reduce out-of-date operations, client synchronization service 156 may load the tree data structures into memory (e.g., random access memory or a cache of high-speed memory) on startup and perform synchronization functions on the tree data structures in memory. Data capacity is limited on persistent storage devices and conservation of these data resources is important. Data capacity is even more limited and expensive for memory and conservation of these data resources is critical.

[204] Depending on the operating system or client application, filenames for content items can be around 1024 bytes in size and may be the largest data component in a node. For example, with a million nodes, the size of just the filenames for the nodes may reach upwards of 1 gigabyte. As noted above, client synchronization service 156 is configured to aid in the synchronization of the server state and the file system state and the local tree, the sync tree, and the remote tree reflect a synchronized state when all three trees are equivalent. Accordingly, there is likely some redundancy in storing the filenames of content items in the nodes.

[205] Various embodiments of the subject technology aim to decrease the amount of memory needed to store filenames for the tree data structures and reduce the size of nodes by reducing the duplication of filenames. Instead of storing the filename in the node, client synchronization service 156 is configured to store filenames for nodes in the tree data structures in a filename array and a reference to the filename in the node.

As a result, the filename is stored once in the filename array and any node with a content item having that filename may access the filename using the reference stored in the node. In some implementations the reference to the filename stored in the node may be an integer value that represents the offset, location, or position of the filename in the filename array.

[206] FIG. 13 shows an illustration of a filename array 1310, in accordance with various embodiments of the subject technology. Filename array 1310 is shown storing filenames “Pictures,” “a.jpg,” and “Documents” (not shown entirely). The “Pictures” filename is shown at location 0, the “a.jpg” filename is shown at location 9, and the “Documents” filename is shown at location 16 in filename array 1310. In filename array 1310, a separator 1314 (e.g., a null character) separates each filename.

[207] Accordingly, to lookup the filename for a node, client synchronization service 156 may simply access the reference to the filename stored in the node, which represents the location of the filename for the node in filename array 1310. Client synchronization service 156 may retrieve the filename of the content item at the location in the filename array 1310 specified by the reference. For example, client synchronization service 156 may begin reading the filename at the location specified in the reference and stop when a separator 1314 is reached.

[208] In some cases, client synchronization service 156 may also need to lookup reference based on a filename. For example, when adding a new node or renaming a node client synchronization service 156 may wish to determine whether the filename of the node already exists in filename array 1310. If the reference is found, the filename exists and may be located based on the reference. Accordingly, client synchronization service 156 may use the reference to the filename and store the reference in the new or renamed node. If the reference is not found, the filename does not exist in filename array 130 and client synchronization service 156 may add the filename to filename array 1310.

[209] The lookup of the reference based on the filename is enabled by the use of a hash index array. FIG. 13 shows an illustration of a hash index array 1350, in accordance with various embodiments of the subject technology. Hash index array 1350 is configured to store references to filenames at various locations. In particular, the reference to a filename is stored at the position in hash index array 1350 based on

the hash of the filename. In one example illustrated in FIG. 13, based on a hash function being used, the hash of filename “Pictures” may equal 4. Accordingly the reference to the filename “Pictures” is stored at position 4 in hash index array 1350. This reference value is 0 which, as noted above, specifies the location in the filename array 1310 of the filename “Pictures.” Similarly, the hash of filename “Documents” may equal 0. Accordingly the reference to the filename “Documents” is stored at position 0 in hash index array 1350. This reference value is 16 which, as noted above, specifies the location in the filename array 1310 of the filename “Documents.”

[210] In some scenarios, a collision may occur where the hash of two filenames yields the same hash value. In the example illustrated in FIG. 13, the hash value of “a.jpg” may also be 4. If a collision occurs, client synchronization service 156 may use the next available position in hash index array 1350. For example, since the reference to the filename “Pictures” is stored at position 4 in hash index array 1350, client synchronization service 156 may look for the next available position in hash index array 1350 and store the reference (e.g., 9) to the filename “a.jpg”

[211] FIG. 14 shows an example method for storing a filename, in accordance with various embodiments of the subject technology. Although the methods and processes described herein may be shown with certain steps and operations in a particular order, additional, fewer, or alternative steps and operations performed in similar or alternative orders, or in parallel, are within the scope of various embodiments unless otherwise stated. The method 1400 may be implemented by a system such as, for example, client synchronization service 156 of FIG. 2, running on a client device.

[212] At operation 1405, the system may detect a modification to a node in a tree data structure. The modification may be, for example, an add of the node to the tree data structure or an edit of the filename of the node. The system may then determine whether or not the filename already exists in the filename array at operation 1410. The system may check to see whether the filename is already in the filename array by querying, using the filename associated with the node, for a reference location of the filename in the filename array. If a reference location is found, there is no need to add the filename to the filename array and at operation 1415, the system may determine the location for the filename in the filename array, which should be the reference returned by the query, and store the location of the filename in the node at operation 1420.

[213] If the query does not return a reference for the filename in the filename array or if the filename is not otherwise in the filename array, the system may append the filename and a separator to the filename array at operation 1425, determine the location of the filename in the filename array at operation 1430, and store the location of the filename in the node at operation 1435.

[214] In order enable the subsequent querying of the reference location based on the filename, the system may further store the location of the filename in a hash index. To determine which position in the hash index to store the location of the filename, the system may compute a hash value of the file name. This hash value may be used to find the position in the hash index to store the location of the filename.

[215] Once the location of the filename is stored in a node, retrieving the filename may simply involve accessing the location of the filename in the node and using the location to lookup the filename in the filename array. The system may start from the location specified in the node and stop when a separator is reached.

[216] FIG. 15 shows an example method for retrieving a location of a filename given the filename, in accordance with various embodiments of the subject technology. Although the methods and processes described herein may be shown with certain steps and operations in a particular order, additional, fewer, or alternative steps and operations performed in similar or alternative orders, or in parallel, are within the scope of various embodiments unless otherwise stated. The method 1500 may be implemented by a system such as, for example, client synchronization service 156 of FIG. 2, running on a client device. As described above, the method may be used to determine whether a filename has already been stored in the filename array.

[217] The locations of the filenames are stored in a hash index or hash index array. Accordingly, at operation 1505, the system may generate a position in the hash index by performing a hash function on the filename. This may be, for example, the name of a new node or the new name for a node being renamed. At operation 1510, the system retrieves, from the position in the hash index, location information of the filename. This location information is for a location in the filename array where the filename is stored.

[218] In some implementations, the system may check to make sure the correct filename is stored in the location at operation 1515. For example, the system may

retrieve, based on the location information, a string from the filename array and compare the string with the filename. If the string and the filename match, the location of the filename is confirmed and, at operation 1520, the location of the filename is stored in the node.

[219] If the string and the filename do not match, the retrieves the location information in the next position in the hash index at operation 1530 and the system may return to operation 1515 to check whether the location information in the next position is accurate. The system may continue until the correct location information is found and, at operation 1520, stored in the node.

Efficiently Identifying Differences Between Trees

[220] As noted above, client synchronization service 156 is configured to identify differences between nodes in a remote tree representing a server state for content items stored by the content management system, a local tree representing the file system state for the corresponding content items stored on the client device, and a sync tree representing a known synced state between the server state and the file system state. Based on these differences, a set of operations may be generated that, if executed, are configured to converge the server state and the file system state towards a synchronized state where the three tree data structures would be identical.

[221] When there are a large number of nodes, it is important to be able to identify differences between the trees efficiently. For example, there may be millions upon millions of nodes in a tree and comparing each node individually can be prohibitive in terms of processing time and resource usage.

[222] Various embodiments of the subject technology relate to providing a more efficient means to identify difference between trees. In particular, client synchronization service 156 is configured to assign values to each node that can be used to compare with nodes in other trees to determine whether the nodes are different. These values may be referred to as diff values. In order to improve the efficiency of identifying differences between trees, each leaf node may be assigned a diff value and diff values for parent nodes may be calculated based on the diff values of their child nodes. Diff values for every level of the tree data structure may be calculated in this way including the root node.

[223] FIGS. 16A and 16B show examples of tree data structures, in accordance with various embodiments. For illustrative purposes, FIGS. 16A and 16B show a sync tree and a local tree. However, the remote tree may perform similarly. In FIGS. 16A and 16B, diff values for the leaf nodes for the sync tree and the local tree are computed using hash functions on the nodes. Each parent node (or node with children) has a diff value computed based on the diff values of their child nodes.

[224] For example, in sync tree 1605 of FIG. 16A, the diff value for node C is 325, the diff value for node D is 742, and the diff value for B, the parent of nodes C and D, is calculated as a function of the diff values of C and D. In other words, the $\text{DiffValue}(B) = f(\text{DiffValue}(C), \text{DiffValue}(D)) = f(325, 742)$. Similarly, the diff value of the root node is a function of the diff values of its child nodes, node A and node B. Put another way, $\text{DiffValue}(/root) = f(\text{DiffValue}(A), \text{DiffValue}(B)) = f(924, 789)$.

[225] In order to identify differences between trees, client synchronization service 156 may compare diff values of corresponding nodes to see whether they are different. If the diff values are the same, there is no difference between trees. If the diff values are different, there is a change in the node or down a path associated with the node. As a result, client synchronization service 156 can look at the child nodes and compare the corresponding nodes in the opposite tree to determine whether they are different and whether there is a difference down a path associated with the child node.

[226] For example, in FIG. 16A, client synchronization service 156 may compare diff values of the root node of sync tree 1605 and the root node of local tree 1610 and determine that the diff values match. Accordingly there is no difference in the children of the root nodes. As a result, client synchronization service 156 may determine that sync tree 1605 and local tree 1610 match and are synchronized without needing to compare each and every node in the trees with the corresponding node in the opposite tree.

[227] In the example illustrated in FIG. 16B, client synchronization service 156 may similarly compare diff values of the root node of sync tree 1605 and the root node of local tree 1610. However, client synchronization service 156 may determine that the diff values of the root node of sync tree 1605 and the root node of local tree 1610 are

different, which indicates that there may be a difference in the descendants of the root nodes.

[228] Accordingly, client synchronization service 156 may move to the next level of the children of the root nodes to compare their diff values. For node A, the diff value of node A in sync tree 1615 and the diff value of node A in local tree 1620 match. Accordingly, there is no difference between the two and the difference detected at the root level is not caused by node A or any path down node A.

[229] Moving to node B, client synchronization service 156 compare their diff values of node B in both trees and discover that the diff value of node A in sync tree 1615 and the diff value of node A in local tree 1620 are different. Accordingly, the difference detected at the root level is caused by node B or a path down node B. Client synchronization service 156 may move to the next level of the children of node B to compare their diff values and discover that the difference detected at the root level was caused by a deletion of node D. As a result, client synchronization service 156 may identify a difference between sync tree 1615 and local tree 1620 without needing to compare each and every node in the trees with the corresponding node in the opposite tree. For example node A may have had many descendant nodes that did not need to be analyzed because the diff values of node A in both trees matched.

[230] A Merkle tree or hash tree mechanism may work in some cases. For example, the diff value for each leaf node may be calculated based on the hash of the leaf node and the diff value of non-leaf nodes may be computed based on a hash of the sum of the child diff values. However, with the Merkle tree or hash tree mechanism has unfavorable performance metrics in certain circumstances. When you add or delete a leaf node, the diff values of all of the ancestor nodes of the leaf node need to have their diff values recomputed and recomputing each of the ancestor nodes requires a listing of all of the children nodes for each ancestor node. This is computationally expensive, especially when the tree data structures are stored in a way where each node is not stored in memory next or near to a sibling node.

[231] Various embodiments of the subject technology address these and other technical shortcomings by, among other things, computing the diff values of the parent nodes differently. The diff value of each leaf node may be determined by calculating a hash of the leaf node. The diff value of each parent node may be

calculated by performing an exclusive-or operation or XOR operation of the hashes of all of its children.

[232] FIG. 17 shows an example of tree data structure, in accordance with various embodiments. In the tree data structure of FIG. 17, the diff values of leaf nodes A, C, and D are calculated by hashing the nodes. The diff value of node B is a XOR function of the hash of its child nodes, node C and node D. The diff value of the root node is a XOR function of the hash of its child nodes, node A and node B. Important qualities of the XOR function include that the XOR function is not order sensitive, meaning that $m \text{ XOR } n$ would equal $n \text{ XOR } m$. Also $m \text{ XOR } m$ equals 0.

[233] When you add or delete a leaf node, the diff values of all of the ancestor nodes of the leaf node need to have their diff values recomputed. However, computing the diff values of the ancestor nodes can be done without requiring a listing of all. For example, when a child node is deleted, the new diff value of the parent may be calculated by performing an XOR operation on the old diff value of the parent and the diff value of the deleted child node. When a child node is added, the new diff value of the parent may be calculated by performing an XOR operation on the old diff value of the parent and the diff value of the new child node.

[234] FIG. 18 shows an example method for retrieving a location of a filename given the filename, in accordance with various embodiments of the subject technology. Although the methods and processes described herein may be shown with certain steps and operations in a particular order, additional, fewer, or alternative steps and operations performed in similar or alternative orders, or in parallel, are within the scope of various embodiments unless otherwise stated. The method 1800 may be implemented by a system such as, for example, client synchronization service 156 of FIG. 2, running on a client device. As described above, the method may be used to determine whether a filename has already been stored in the filename array.

[235] At operation 1805, client synchronization service 156 may add or remove a node from a tree data structure such as the remote tree, the sync tree, or the local tree. If the node is deleted, the node should have a previously calculated diff value associated with it. If the node is added, client synchronization service 156 may calculate a diff value for the new node by, for example, hashing the node. Because

there is a change to the tree, the diff values for the ancestors of the node must be updated.

[236] At operations 1810, client synchronization service 156 may calculate a new diff value for the parent of the node based on the current diff value of the parent node and the diff value of the node. At operation 1815, the new diff value for the parent is stored in the parent node.

[237] At operation 1820, client synchronization service 156 determines whether the parent node itself has a parent. In other words, where the parent node a root node or whether there additional ancestors to compute diff values for. If there is an additional parent, the process may return to operation 1810 where the parent of the parent node will have a new diff value calculated and stored. If there is not an additional parent and the parent node is a root node, the process may stop at operation 1825.

[238] Once the root node is reached the tree data structure is ready to be compared with other tree data structures in order to identify differences. As noted above, client synchronization service 156 may generate a set of operations based on these differences that if executed, are configured to converge the server state and the file system state towards a synchronized state where the three tree data structures would be identical.

[239] FIG. 19 shows an example of computing system 1900, which can be for example any computing device making up client device 150, content management system 110 or any component thereof in which the components of the system are in communication with each other using connection 1905. Connection 1905 can be a physical connection via a bus, or a direct connection into processor 1910, such as in a chipset architecture. Connection 1905 can also be a virtual connection, networked connection, or logical connection.

[240] In some embodiments computing system 1900 is a distributed system in which the functions described in this disclosure can be distributed within a datacenter, multiple datacenters, a peer network, etc. In some embodiments, one or more of the described system components represents many such components each performing some or all of the function for which the component is described. In some embodiments, the components can be physical or virtual devices.

[241] Example system 1900 includes at least one processing unit (CPU or processor) 1910 and connection 1905 that couples various system components including system memory 1915, such as read only memory (ROM) 1920 and random access memory (RAM) 1925 to processor 1910. Computing system 1900 can include a cache of high-speed memory 1912 connected directly with, in close proximity to, or integrated as part of processor 1910.

[242] Processor 1910 can include any general purpose processor and a hardware service or software service, such as services 1932, 1934, and 1936 stored in storage device 1930, configured to control processor 1910 as well as a special-purpose processor where software instructions are incorporated into the actual processor design. Processor 1910 may essentially be a completely self-contained computing system, containing multiple cores or processors, a bus, memory controller, cache, etc. A multi-core processor may be symmetric or asymmetric.

[243] To enable user interaction, computing system 1900 includes an input device 1945, which can represent any number of input mechanisms, such as a microphone for speech, a touch-sensitive screen for gesture or graphical input, keyboard, mouse, motion input, speech, etc. Computing system 1900 can also include output device 1935, which can be one or more of a number of output mechanisms known to those of skill in the art. In some instances, multimodal systems can enable a user to provide multiple types of input/output to communicate with computing system 1900. Computing system 1900 can include communications interface 1940, which can generally govern and manage the user input and system output. There is no restriction on operating on any particular hardware arrangement and therefore the basic features here may easily be substituted for improved hardware or firmware arrangements as they are developed.

[244] Storage device 1930 can be a non-volatile memory device and can be a hard disk or other types of computer readable media which can store data that are accessible by a computer, such as magnetic cassettes, flash memory cards, solid state memory devices, digital versatile disks, cartridges, random access memories (RAMs), read only memory (ROM), and/or some combination of these devices.

[245] The storage device 1930 can include software services, servers, services, etc., that when the code that defines such software is executed by the processor 1910, it

causes the system to perform a function. In some embodiments, a hardware service that performs a particular function can include the software component stored in a computer-readable medium in connection with the necessary hardware components, such as processor 1910, connection 1905, output device 1935, etc., to carry out the function.

[246] For clarity of explanation, in some instances the present technology may be presented as including individual functional blocks including functional blocks comprising devices, device components, steps or routines in a method embodied in software, or combinations of hardware and software.

[247] Any of the steps, operations, functions, or processes described herein may be performed or implemented by a combination of hardware and software services or services, alone or in combination with other devices. In some embodiments, a service can be software that resides in memory of a client device and/or one or more servers of a content management system and perform one or more functions when a processor executes the software associated with the service. In some embodiments, a service is a program, or a collection of programs that carry out a specific function. In some embodiments, a service can be considered a server. The memory can be a non-transitory computer-readable medium.

[248] In some embodiments the computer-readable storage devices, mediums, and memories can include a cable or wireless signal containing a bit stream and the like. However, when mentioned, non-transitory computer-readable storage media expressly exclude media such as energy, carrier signals, electromagnetic waves, and signals per se.

[249] Methods according to the above-described examples can be implemented using computer-executable instructions that are stored or otherwise available from computer readable media. Such instructions can comprise, for example, instructions and data which cause or otherwise configure a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Portions of computer resources used can be accessible over a network. The computer executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, firmware, or source code. Examples of computer-readable media that may be used to store instructions,

information used, and/or information created during methods according to described examples include magnetic or optical disks, solid state memory devices, flash memory, USB devices provided with non-volatile memory, networked storage devices, and so on.

[250] Devices implementing methods according to these disclosures can comprise hardware, firmware and/or software, and can take any of a variety of form factors. Typical examples of such form factors include servers, laptops, smart phones, small form factor personal computers, personal digital assistants, and so on. Functionality described herein also can be embodied in peripherals or add-in cards. Such functionality can also be implemented on a circuit board among different chips or different processes executing in a single device, by way of further example.

[251] The instructions, media for conveying such instructions, computing resources for executing them, and other structures for supporting such computing resources are means for providing the functions described in these disclosures.

[252] Although a variety of examples and other information was used to explain aspects within the scope of the appended claims, no limitation of the claims should be implied based on particular features or arrangements in such examples, as one of ordinary skill would be able to use these examples to derive a wide variety of implementations. Further and although some subject matter may have been described in language specific to examples of structural features and/or method steps, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to these described features or acts. For example, such functionality can be distributed differently or performed in components other than those identified herein. Rather, the described features and steps are disclosed as examples of components of systems and methods within the scope of the appended claims.

CLAUSES

[253] A computer-implemented method comprising: comparing a remote tree and a sync tree to identify a server difference between the remote tree and the sync tree, wherein the remote tree represents a server state for content items stored by a content management system during a first time period, and wherein the sync tree represents a synced state during a previous time period when the remote tree matched a local tree, wherein the local tree represents a local state for content items stored by the client

device; determining, based on the server difference, that the server state and the file system state are out of sync; generating, based on the server difference, a client set of operations for the content items stored on the client device, wherein the client set of operations is configured to operate on the content items stored on the client device to converge the server state and the file system state; and providing the client set of operations for execution on the client device.

[254] A non-transitory computer readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to: identify a client difference between a local tree and a sync tree, wherein the local tree represents a file system state for content items stored on the computing system during a first time period, and wherein the sync tree represents a synced state during a previous time period when the local tree matched a remote tree, wherein the remote tree represents a server state for content items stored by a content management system; determine, based on a client difference between the local tree and the sync tree, that a server state and a file system state are out of sync; generate, based on the client difference, a server set of operations for the content items stored by the content management system, wherein the server set of operations is configured to operate on the content items stored by the content management system to converge the server state and the file system state; and transmit the server set of operations to the content management system for execution.

[255] A system comprising: a processor; and a non-transitory computer-readable medium storing instructions that, when executed by the processor, cause the processor to: generate a comparison between a sync tree and at least one of a remote tree and a local tree, wherein the remote tree represents a server state during a first time period and the local tree represents a file system state during the first time period, and wherein the sync tree represents a synced state between the server state and the file system state during a previous time period; identify, based on the comparison, at least one difference between a sync tree and at least one of a remote tree and a local tree; generate a set of operations based on the at least one difference, wherein the set of operations is configured to converge the server state and the file system state; and manage execution of the set of operations.

[256] A computer-implemented method for synchronizing content items between a user account on a content management system and a client device authorized to access

the user account, the method comprising: obtaining a set of tree data structures including a remote tree representing a server state of content items associated with the user account on the content management system during a current time period, a local tree representing a file system state of content items associated with the user account on the client device during the current time period, and a sync tree representing a known sync state during a previous time period where the local tree and the remote tree were identical, wherein the sync tree matches the local tree and the remote tree at the previous time period; determining that the user account on the content management system includes at least one modification not synchronized to the client device by comparing the remote tree and the sync tree; and determining that the client device includes at least one modification not synchronized to the content management system by comparing the sync tree and the local tree.

[257] A non-transitory computer readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to: obtain a set of tree data structures including a remote tree representing a server state of content items associated with a user account on a content management system during a current time period, a local tree representing a file system state of content items associated with the user account on a client device during the current time period, and a sync tree representing a known sync state during a previous time period where the local tree and the remote tree were identical, wherein the sync tree matches the local tree and the remote tree at the previous time period; and determine that the user account on the content management system includes at least one modification not synchronized to the client device by comparing the remote tree and the sync tree.

[258] A system comprising: a processor; and a non-transitory computer-readable medium storing instructions that, when executed by the processor, cause the processor to: obtain a set of tree data structures including a remote tree representing a server state of content items associated with a user account on a content management system during a current time period, a local tree representing a file system state of content items associated with the user account on a client device during the current time period, and a sync tree representing a known sync state during a previous time period where the local tree and the remote tree were identical, wherein the sync tree matches the local tree and the remote tree at the previous time period; and determine that the

client device includes at least one modification not synchronized to the content management system by comparing the sync tree and the local tree.

[259] A computer-implemented method for synchronizing modifications to a user account on a content management system to a client device authorized to access the user account, the method comprising: computing a difference between a remote tree data structure representing a server state for content items associated with the user account on the content management system and a sync tree data structure representing a known synchronization state between the content management system and the client device; and generating, based on the difference, a set of operations that when performed on the client device update the content items stored on the client device to converge a file system state on the client device and the server state.

[260] A non-transitory computer readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to: compute a difference between a remote tree data structure representing a server state for content items associated with an account on a content management system and a sync tree data structure representing a known synchronization state between the content management system and the computing system; and generate, based on the difference, a set of operations that when performed on the computing system update the content items stored on the client device to converge a file system state on the computing system and the server state.

[261] A system comprising: a processor; and a non-transitory computer-readable medium storing instructions that, when executed by the processor, cause the processor to: compute a difference between a remote tree data structure representing a server state for content items associated with a user account on a content management system and a sync tree data structure representing a known synchronization state between the content management system and the system; and generate, based on the difference, a set of operations that when performed on the system, update the content items stored on the client device to converge a file system state on the system and the server state.

[262] A computer-implemented method for synchronizing modifications to content items on a client device to a user account on a content management system, the method comprising: computing a difference between a local tree data structure

representing a file system state for content items associated with the user account on the client device and a sync tree data structure representing a known synchronization state between the content management system and the client device; and generating, based on the difference, a set of operations that when performed update the content items stored on the content management system to converge a server state for content items associated the user account on the content management system and the file system state.

[263] A non-transitory computer readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to: compute a difference between a local tree data structure representing a file system state for content items associated with a user account on the computing system and a sync tree data structure representing a known synchronization state between a content management system and the computing system; and generate, based on the difference, a set of operations that when performed update the content items stored on the content management system to converge a server state for content items associated the user account on the content management system and the file system state.

[264] A system comprising: a processor; and a non-transitory computer-readable medium storing instructions that, when executed by the processor, cause the processor to: compute a difference between a local tree data structure representing a file system state for content items associated with a user account on the system and a sync tree data structure representing a known synchronization state between a content management system and the system; and generate, based on the difference, a set of operations that when performed update the content items stored on the content management system to converge a server state for content items associated the user account on the content management system and the file system state.

CLAIMS

What is claimed is:

1. A system comprising:
 - a processor; and
 - a computer-readable medium storing instructions that, when executed by the processor, cause the processor to:
 - generate a comparison between a sync tree and at least one of a remote tree and a local tree, wherein the remote tree represents a server state during a first time period and the local tree represents a file system state during the first time period, and wherein the sync tree represents a synced state between the server state and the file system state during a previous time period;
 - identify, based on the comparison, at least one difference between a sync tree and at least one of a remote tree and a local tree;
 - generate a set of operations based on the at least one difference, wherein the set of operations is configured to converge the server state and the file system state; and
 - manage execution of the set of operations.
2. The system of claim 1, wherein the instructions further cause the processor to:
 - receive server modification data for content items stored by a content management service; and
 - update, based on the server modification data, the remote tree.
3. The system of claim 1, wherein the instructions further cause the processor to:
 - receive client modification data for content items stored by the system; and update, based on the client modification data, the local tree.
4. The system of any one of the preceding claims, wherein the instructions further cause the processor to:
 - identify, for each operation in the set of operations, a task associated with the operation; and
 - perform the task.

5. The system of any one of the preceding claims, wherein the remote tree, local tree, and the sync tree are stored on the non-transitory computer-readable media of the system.
6. The system of claim 3, wherein the server state and the file system state being out of sync indicates that an action performed on the content items stored by the content management system is not reflected in the content items stored on the client device.
7. The system of claim 2, wherein the server state and the file system state being out of sync indicates that an action performed on the content items stored on the computing system is not reflected in the content items stored by the content management system.
8. A computer readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to:
 - compare a remote tree and a sync tree , wherein the remote tree represents a server state during a first time period, and wherein the sync tree represents a synced state between the server state and a file system state for content items stored on the computing system during a previous time period;
 - determine, based on the comparison, at least one difference between the remote tree and the sync tree, the at least one difference indicating that the server state and the file system state are out of sync;
 - generate, based on the determined at least one difference, a set of operations for the content items stored on the computing system, wherein the set of operations is configured to operate on the content items stored on the computing system to converge the server state and the file system state.
9. A computer readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to:
 - obtain a set of tree data structures including a remote tree representing a server state of content items associated with a user account on a content management system during a current time period, a local tree representing a file system state of content items associated with the user account on a client device during the current time period, and a sync tree representing a known sync state during a previous time period where the local tree and the remote tree were identical, wherein the sync tree matches the local tree and the remote tree at

the previous time period; and between the content management system and the client device, wherein at the known sync state the content items associated with the user account on the content management system are synchronized with the content items associated with the user account on the client device; and

determine that the user account on the content management system includes at least one modification not synchronized to the client device by comparing the remote tree and the sync tree.

10. The computer readable medium of claim 9, wherein instructions further cause the computing system to generate a set of operations configured to operate on the content items stored on the client device to converge the server state and the file system state.

11. The computer readable medium of any one of claims 9 or 10, wherein instructions further cause the computing system to determine that the client device includes at least one modification not synchronized to the content management system by comparing the sync tree and the local tree.

12. The computer readable medium of claim 11, wherein instructions further cause the computing system to generate a set of operations configured to operate on the content items stored by the content management system to converge the server state and the file system state.

13. The computer readable medium of any one of claims 9-12, wherein the set of tree data structures is stored on the client device.

14. A computer-implemented method for synchronizing modifications to a user account on a content management system to a client device authorized to access the user account, the method comprising:

computing a difference between a remote tree data structure representing a server state for content items associated with the user account on the content management system and a sync tree data structure representing a known synchronization state between the content management system and the client device; and

generating, based on the difference, a set of operations that when performed on the client device update the content items stored on the client device to converge a file system state on the client device and the server state.

15. The computer- implemented method of claim 14, further comprising executing the set of operations.

16. The computer-implemented method of any one of claims 14 and 15, wherein the remote tree data structure and the sync tree data structure are stored on the client device.

17. The computer-implemented method of any one of claims 14-16, further comprising:
receiving, from the client device, client modification data for content items stored on the client device; and
updating, based on the client modification data, the local tree.

18. The computer- implemented method of claim 17, further comprising updating, based on the client modification data, the sync tree.

19. The computer-implemented method of any one of claims 14-18, further comprising:
computing a second difference between a local tree data structure representing a file system state for content items associated with the user account on the client device and the sync tree data structure representing a known synchronization state between the content management system and the client device;

generating, based on the difference, a second set of operations that when performed update the content items stored on the content management system to converge a server state for content items associated the user account on the content management system and the file system state; and

managing execution of the second set of operations.

20. The computer-implemented method of claim 19, wherein the local tree data structure is stored on the client device.

21. The computer-implemented method of claim 19, further comprising:

receiving, from the content management system, server modification data for content items stored by the content management system; and

updating, based on the server modification data, the remote tree.

22. The computer-implemented method of claim 21, further comprising updating, based on the server modification data, the sync tree.

23. A computer- implemented method for synchronizing modifications to content items on a client device to a user account on a content management system, the method comprising:
computing a difference between a local tree data structure representing a file system state for content items associated with the user account on the client device and a sync tree data structure representing a known synchronization state between the content management system and the client device; and

generating, based on the difference, a set of operations that when performed update the content items stored on the content management system to converge a server state for content items associated the user account on the content management system and the file system state.

24. The computer- implemented method of claim 23, further comprising executing the set of operations.

25. The computer- implemented method of any one of claims 23-24, wherein the local tree data structure and the sync tree data structure are stored on the client device.

26. The computer-implemented method of any one of claims 23-25, further comprising:
receiving, from the client device, client modification data for content items stored on the client device; and

updating, based on the client modification data, the local tree.

27. The computer-implemented method of claim 26, further comprising updating, based on the client modification data, the sync tree.

28. The computer-implemented method of any one of claims 23-27, further comprising:

receiving, from the content management system, server modification data for content items stored by the content management system; and

updating, based on the server modification data, a remote tree.

29. The computer-implemented method of claim 28, further comprising updating, based on the server modification data, the sync tree.

30. The computer-implemented method of any one of claims 23-29, further comprising: computing a second difference between a remote tree data structure representing the server state for content items associated with the user account on the content management system and the sync tree data structure representing the known synchronization state between the content management system and the client device; and

generating, based on the second difference, a second set of operations that when performed on the client device update the content items stored on the client device to converge a file system state on the client device and the server state.

31. The computer-implemented method of claim 30, further comprising managing execution of the second set of operations.

32. The computer-implemented method of claim 30, wherein the remote tree data structure is stored on the client device.


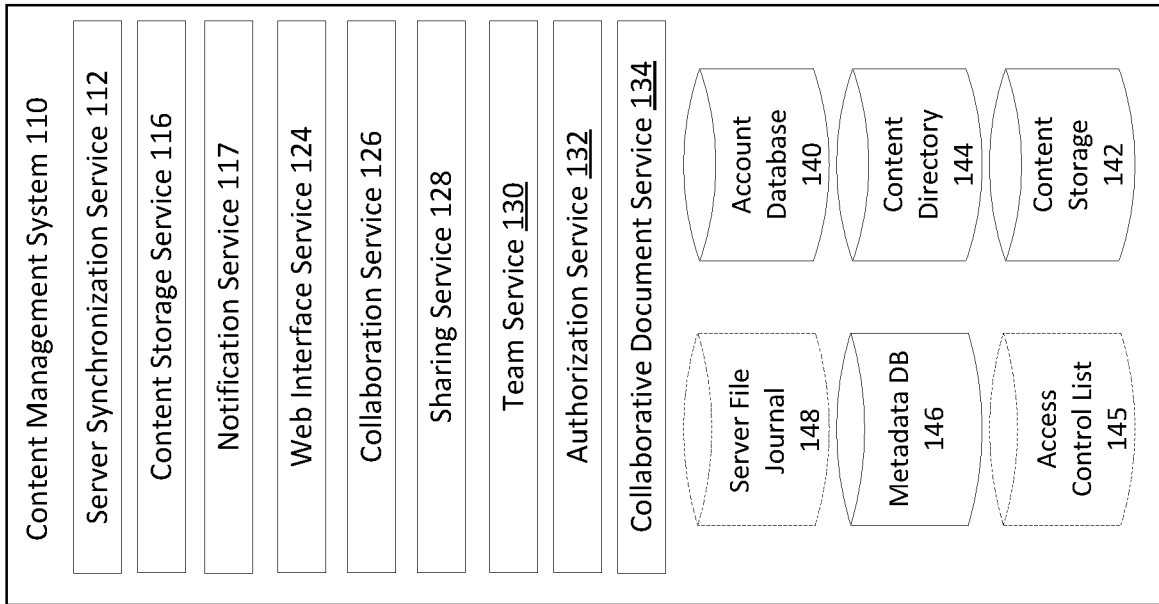


FIG. 1




FIG. 2




FIG. 3


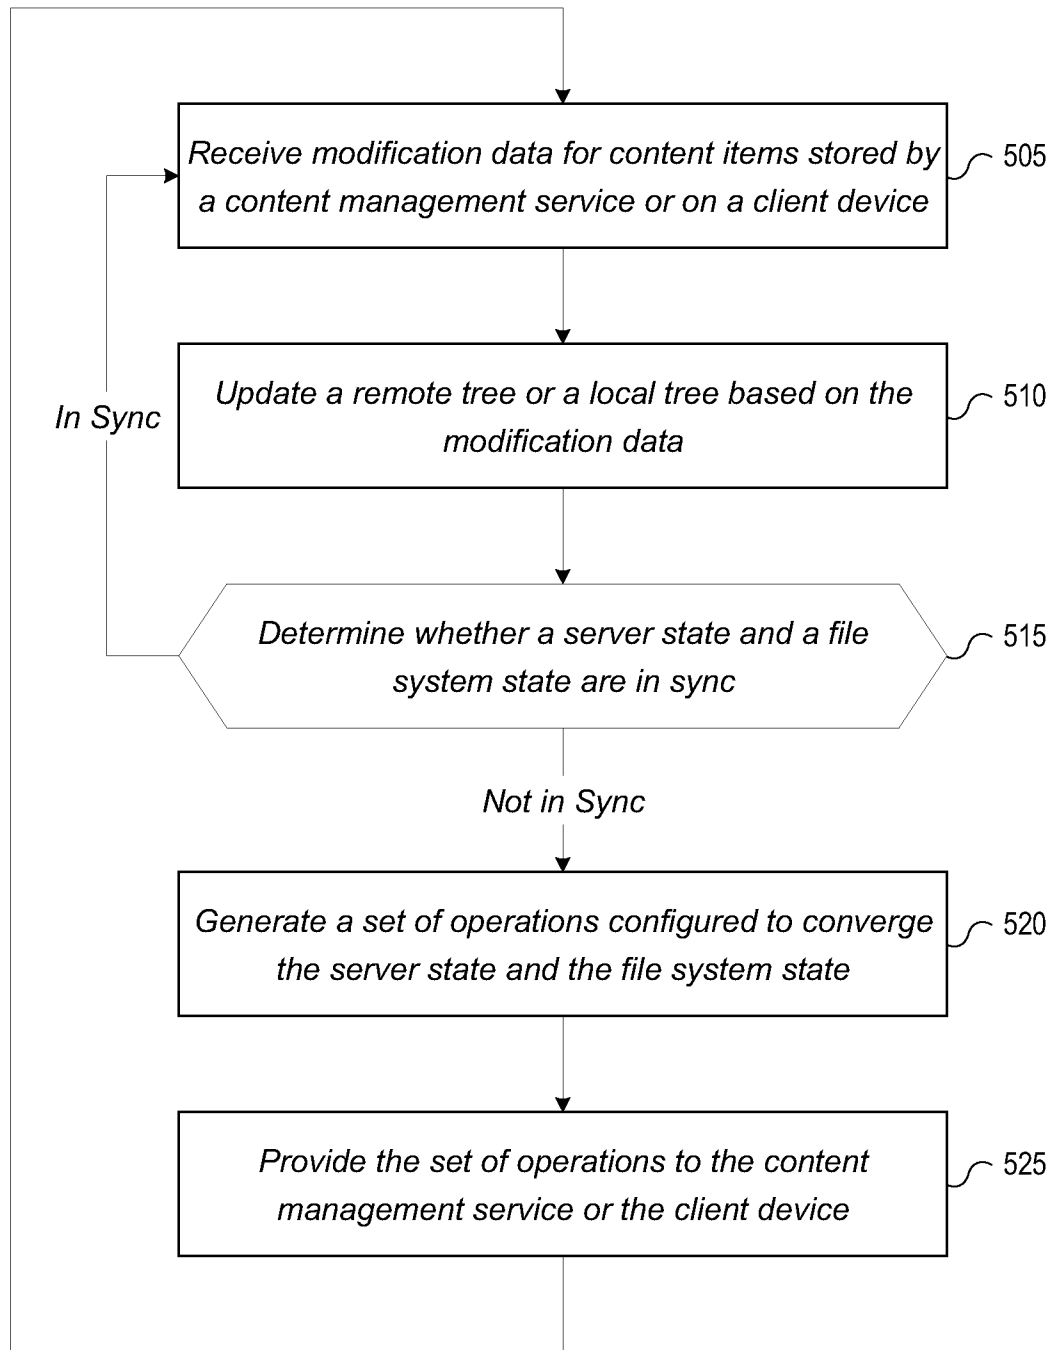
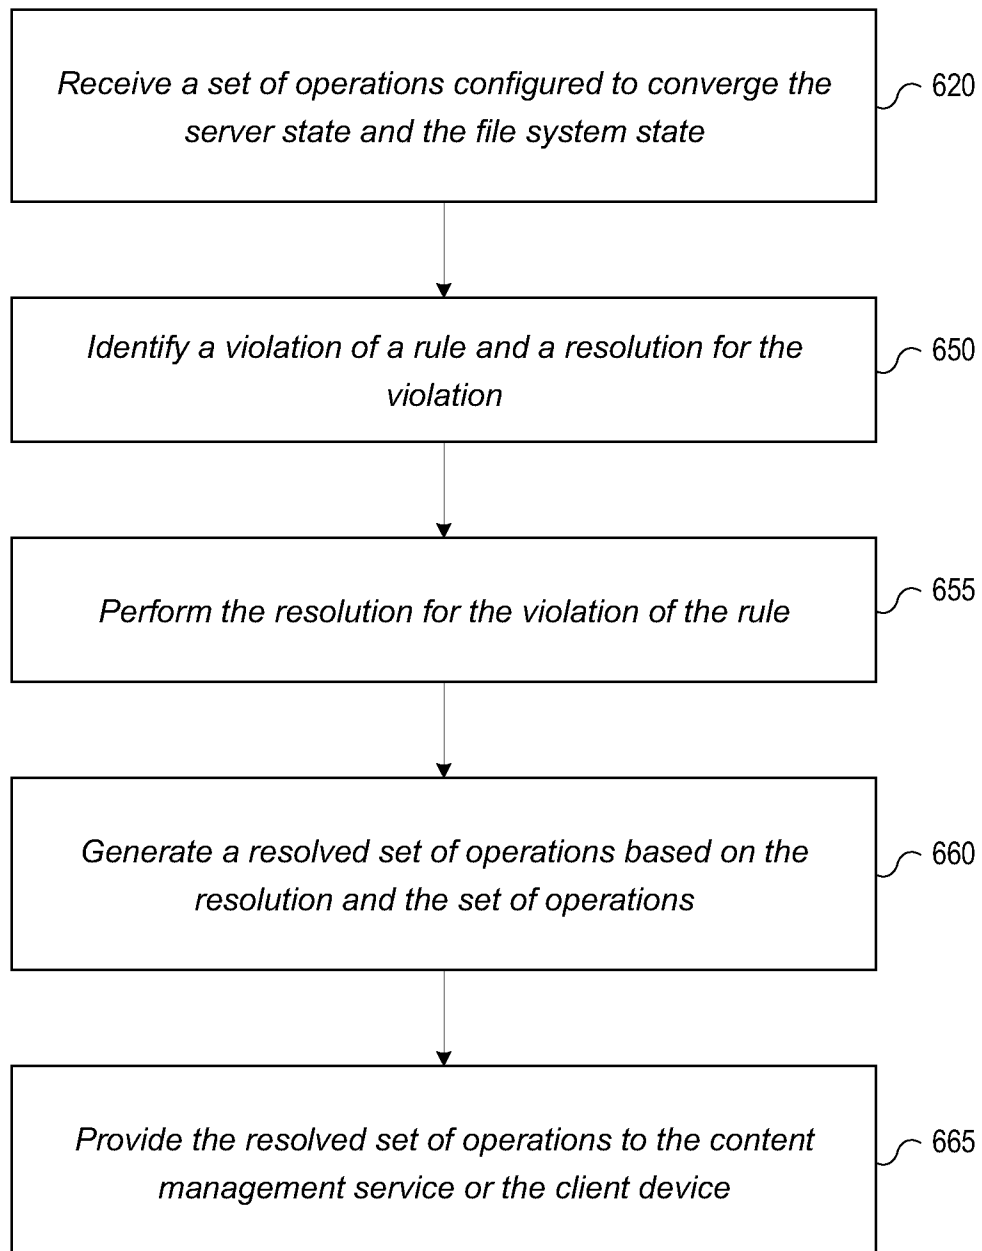


FIG. 4



500

FIG. 5



600
FIG. 6




FIG. 7

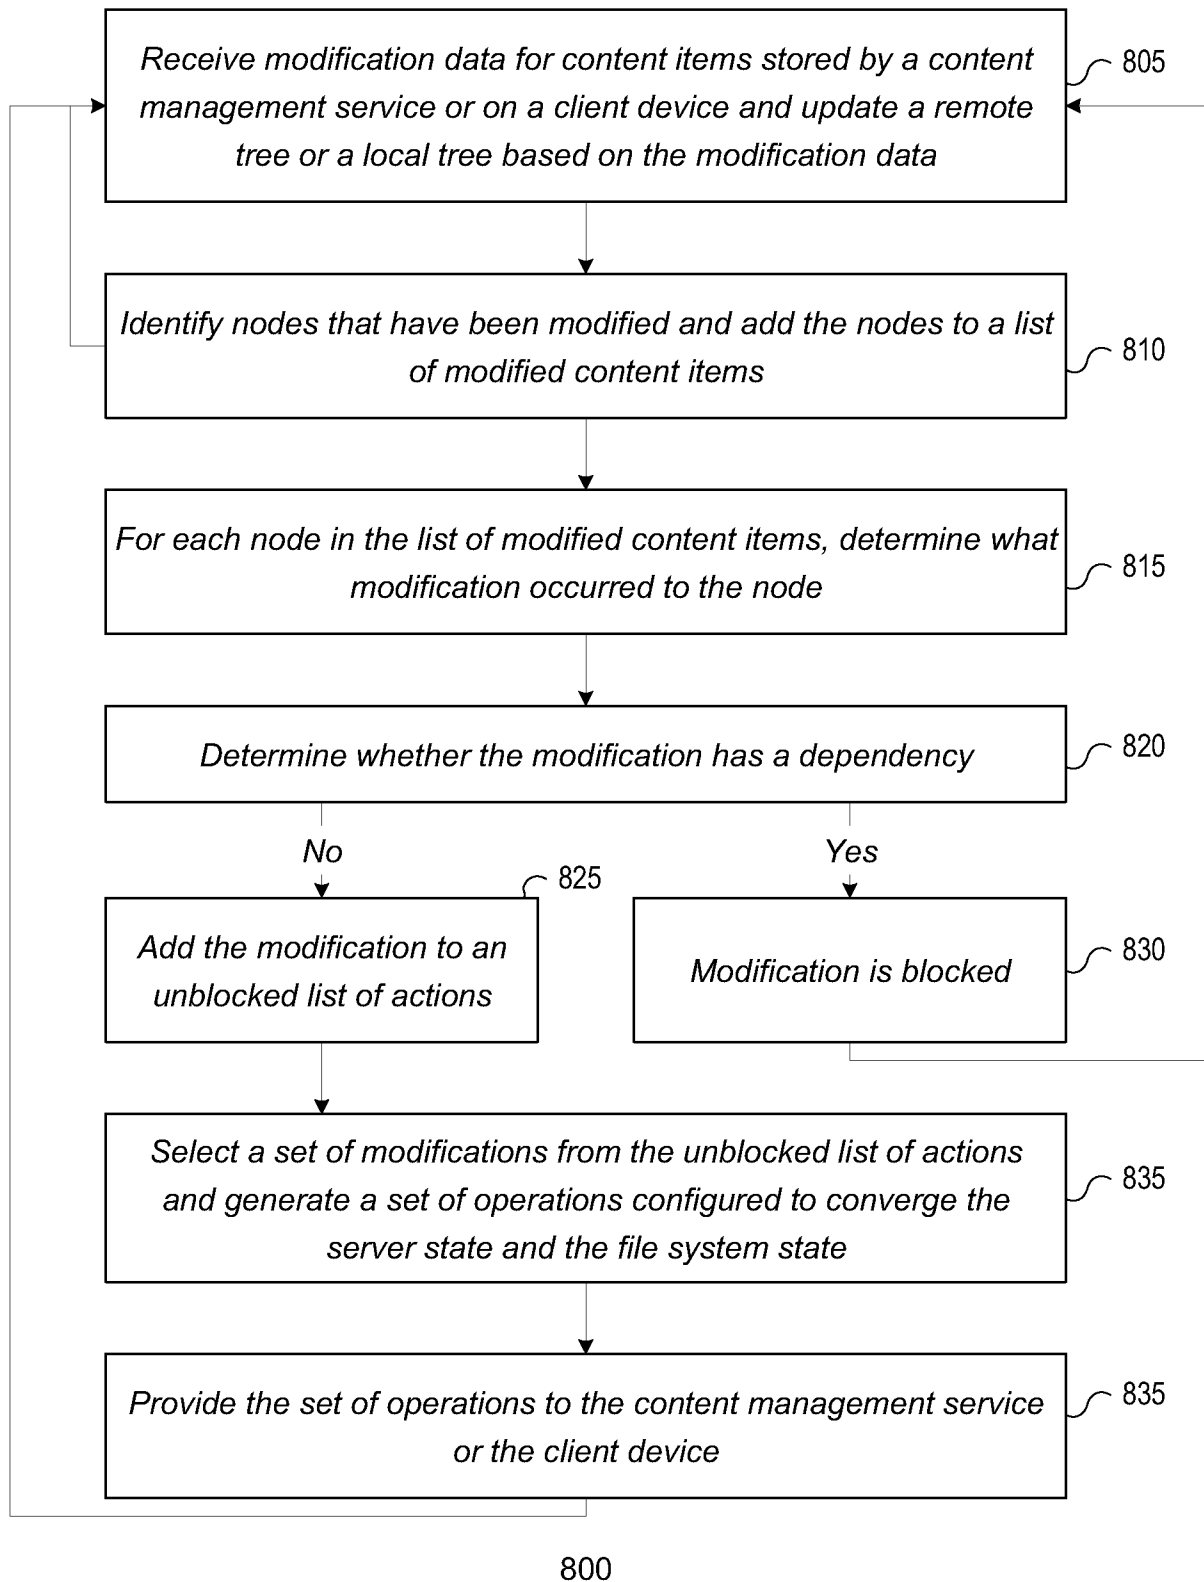


FIG. 8




FIG. 9

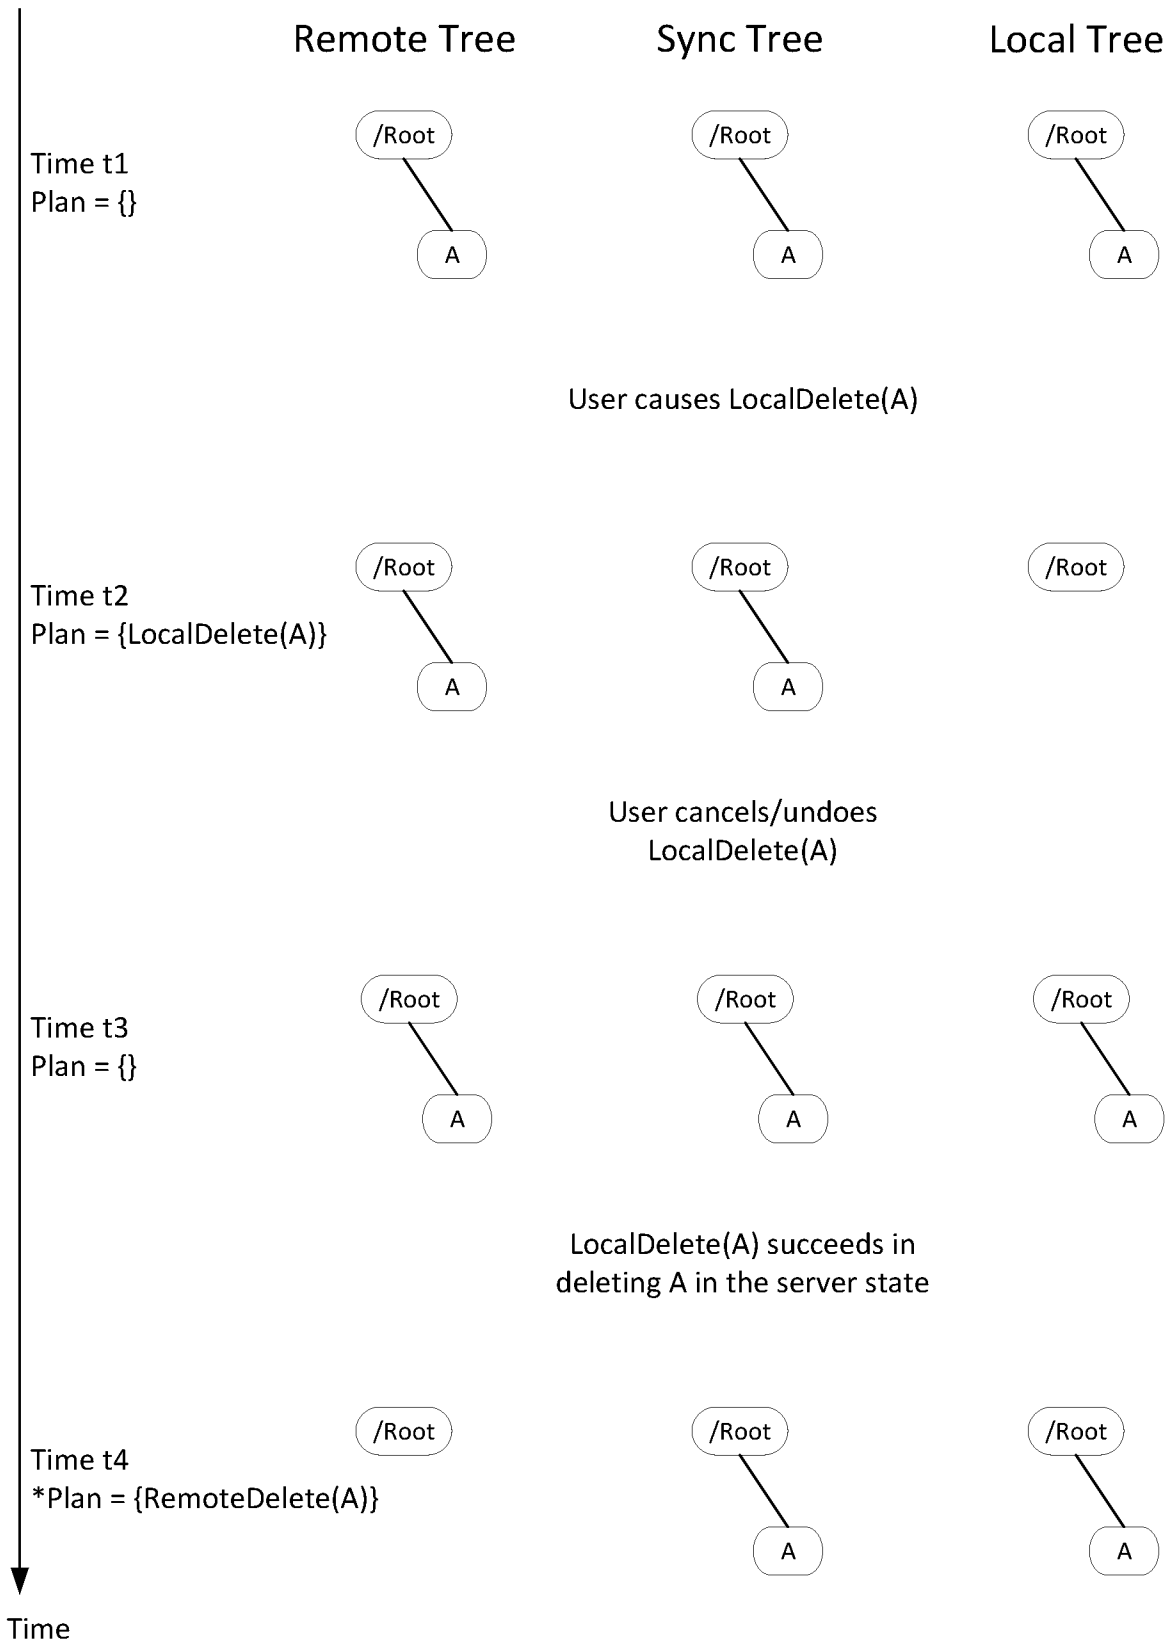


FIG. 10




FIG. 11

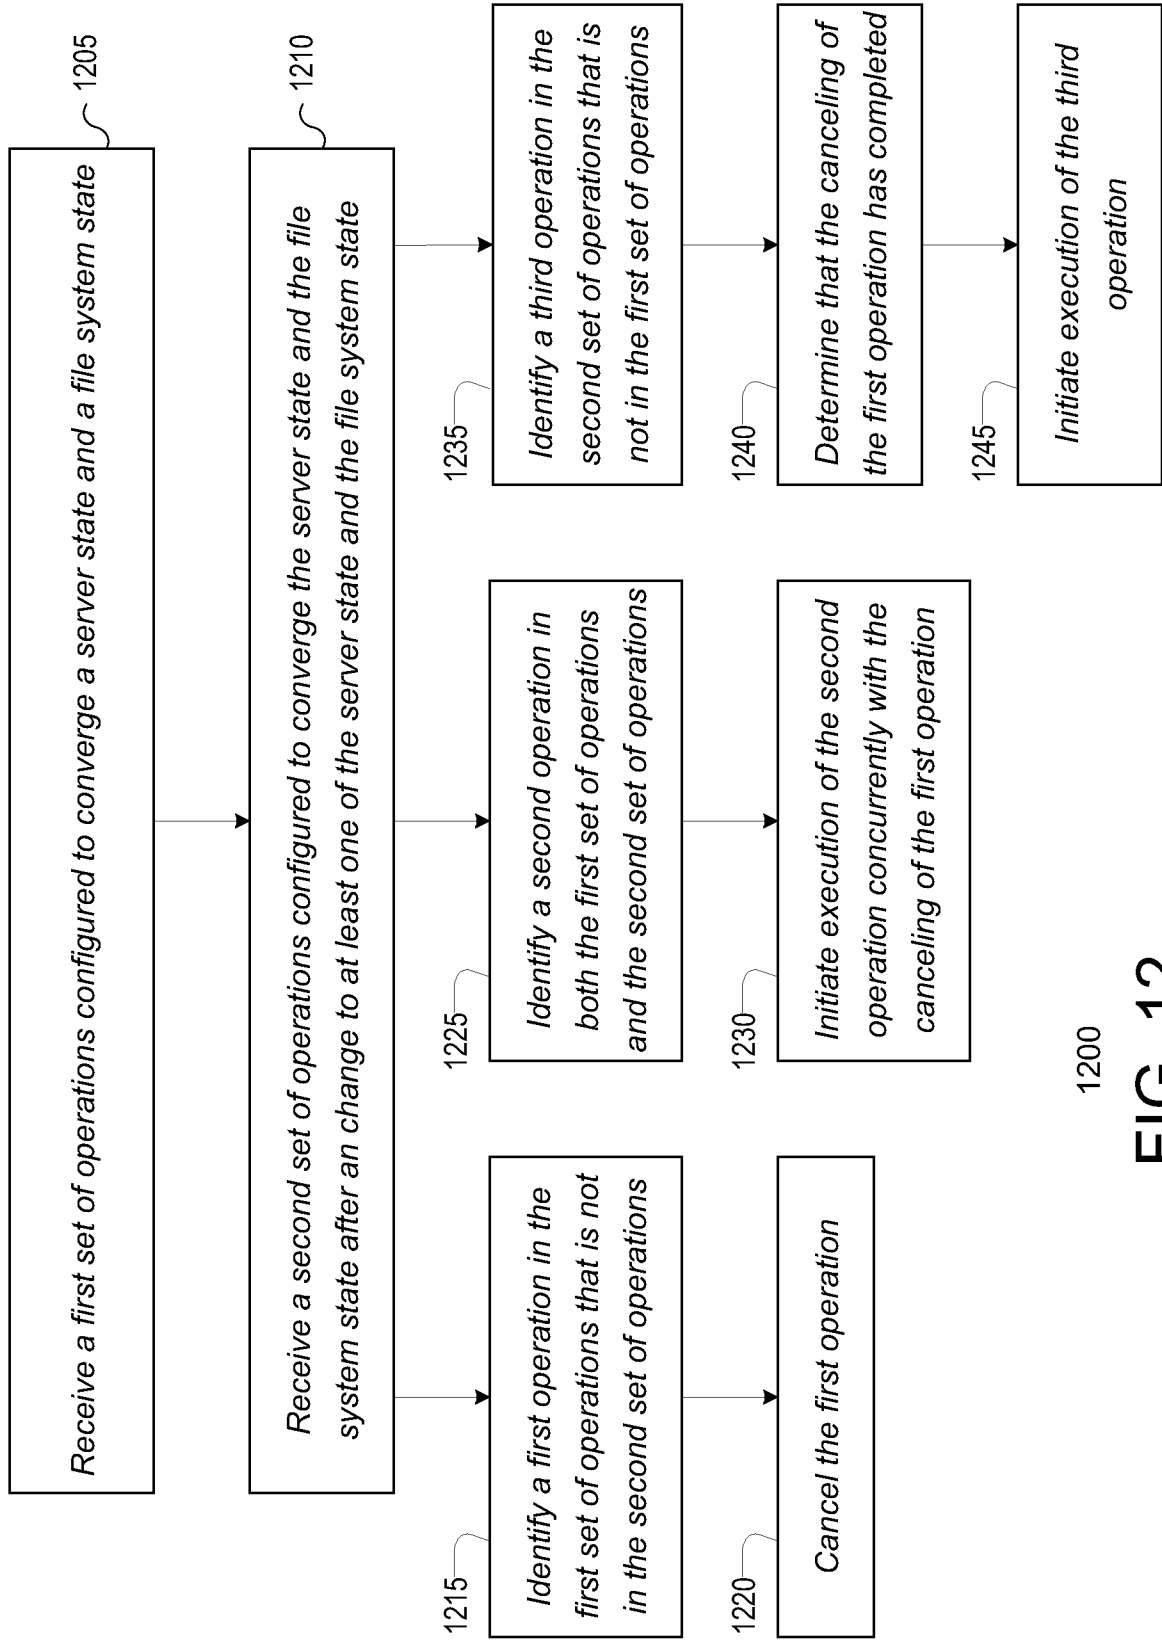


FIG. 12


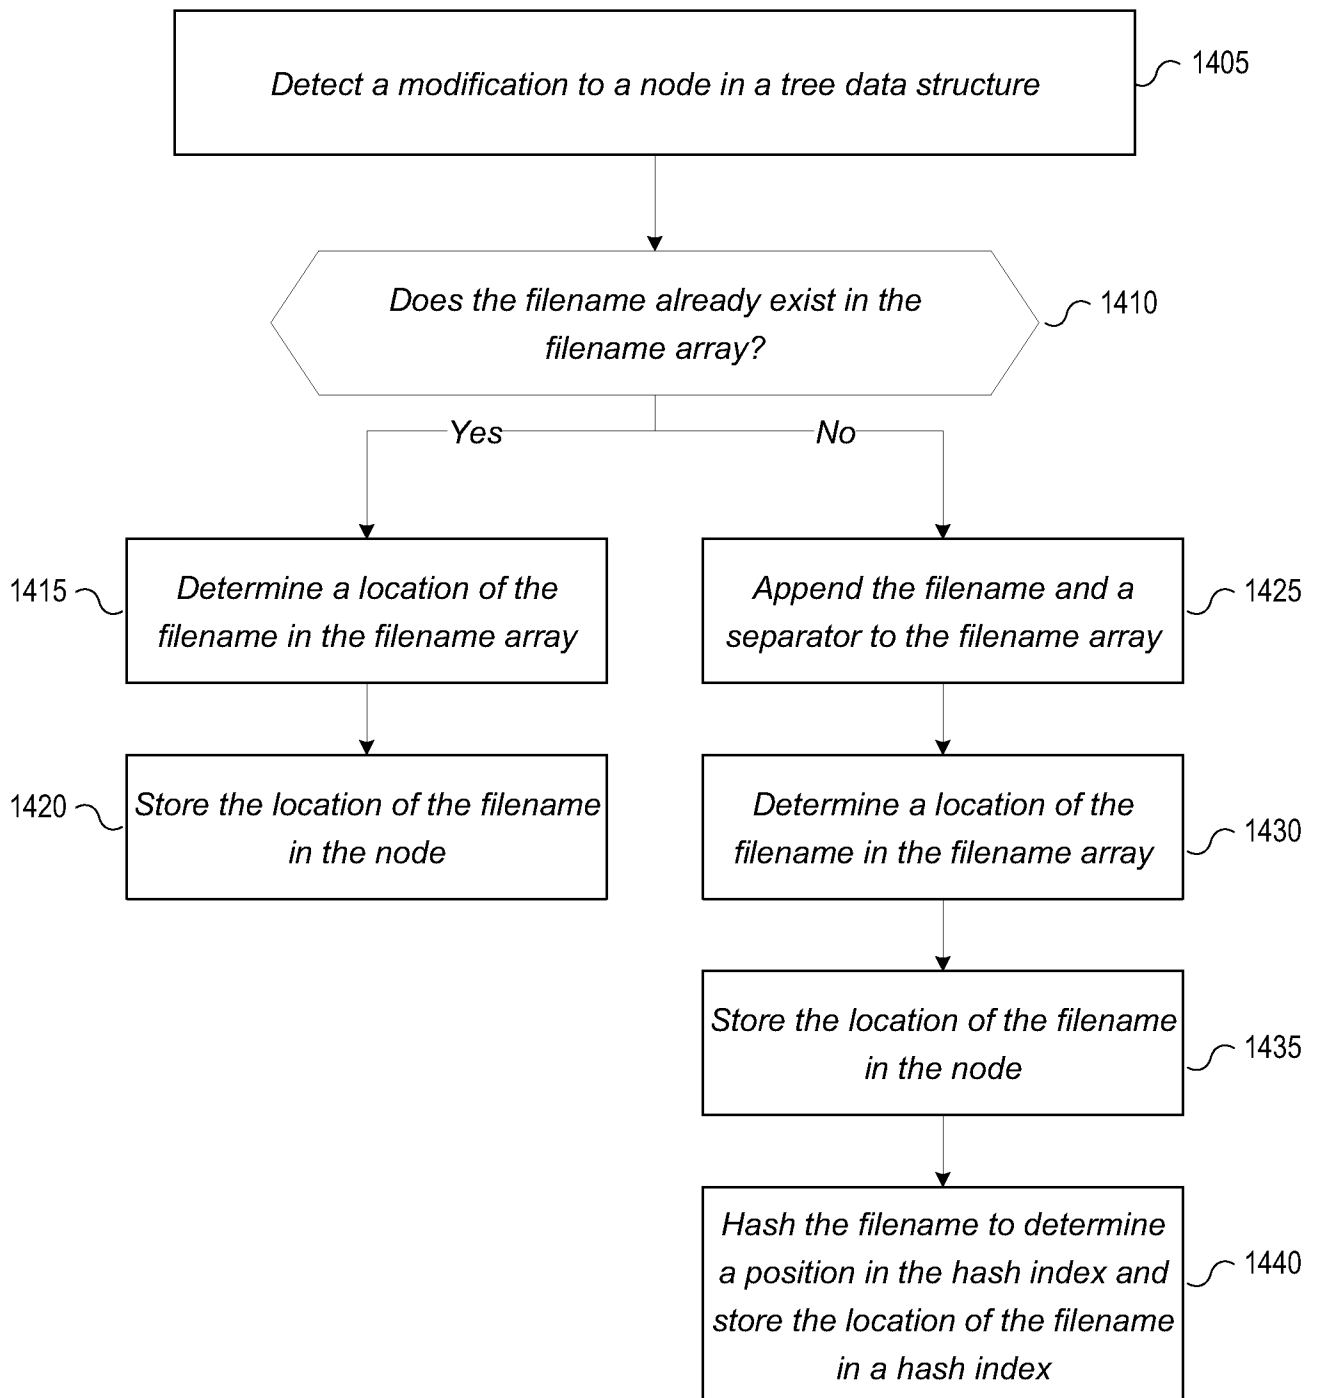
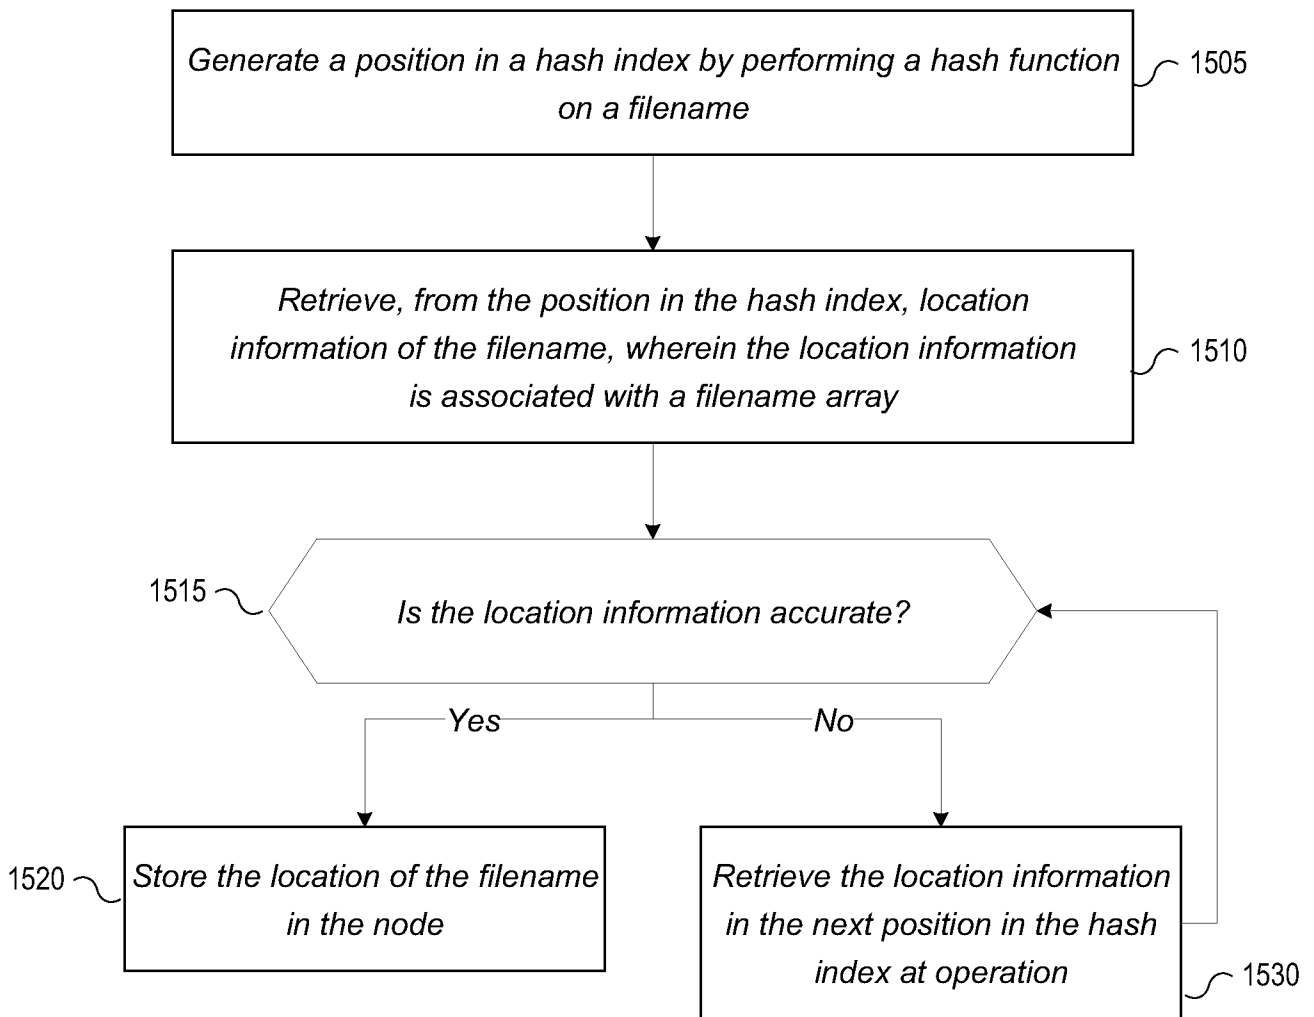


FIG. 13



1400

FIG. 14



1500

FIG. 15




FIG. 16A




FIG. 16B

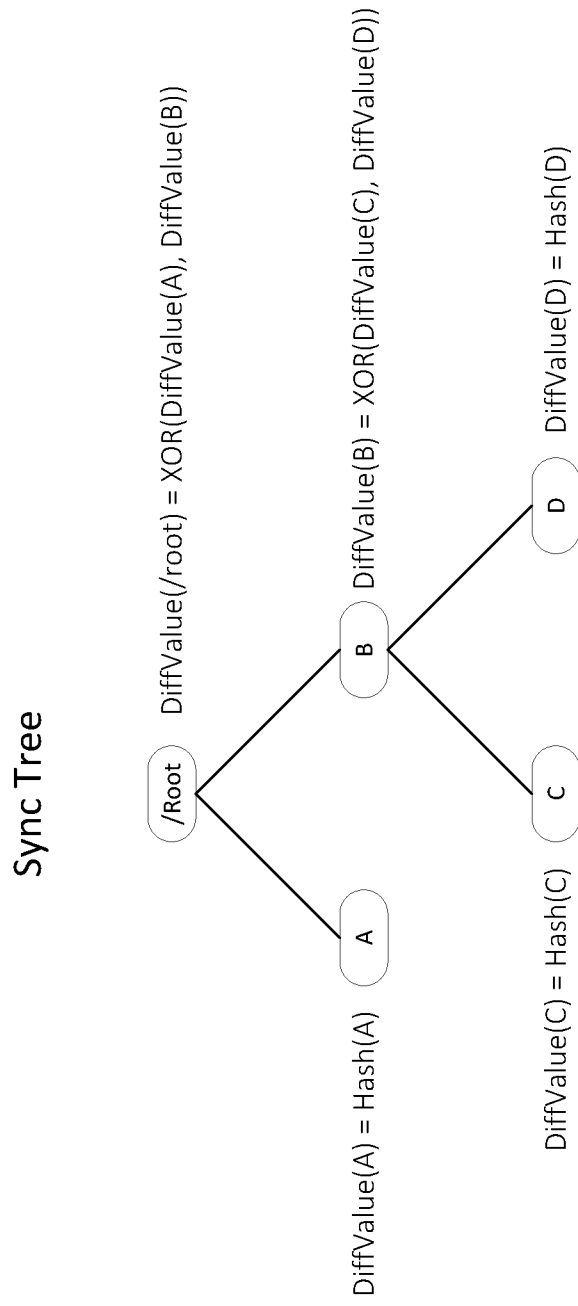
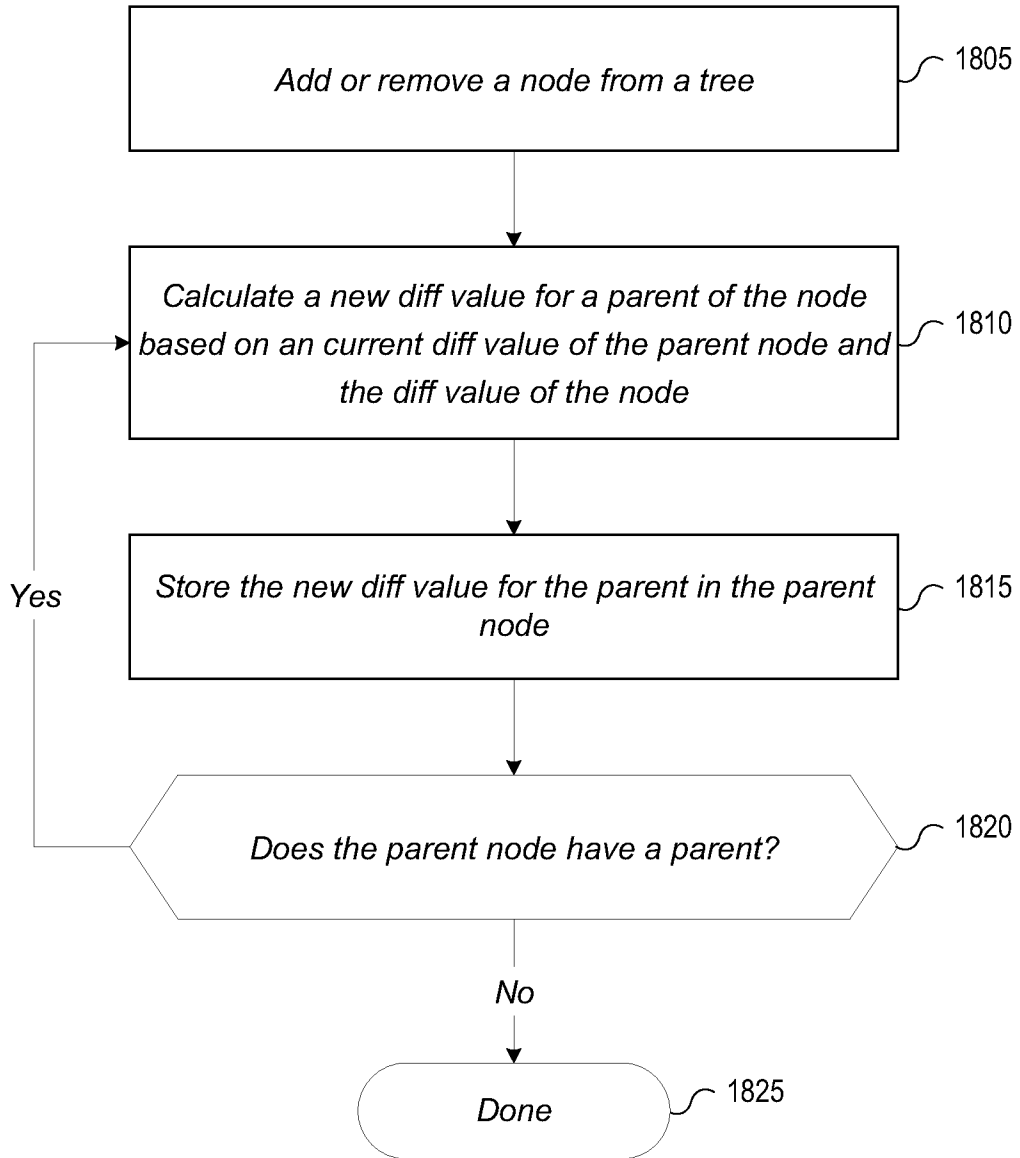


FIG. 17



1800

FIG. 18




FIG. 19