

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
14 May 2009 (14.05.2009)

PCT

(10) International Publication Number
WO 2009/061638 A1(51) International Patent Classification:
G06F 17/21 (2006.01) *G06F 17/24* (2006.01)
G06F 17/00 (2006.01)(21) International Application Number:
PCT/US2008/081456

(22) International Filing Date: 28 October 2008 (28.10.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/938,082 9 November 2007 (09.11.2007) US

(71) Applicant (for all designated States except US): MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).


(72) Inventors: BAILOR, Jonathan Beckett; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). BERNSTEIN,

Ethan Joseph; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). KNIGHT, Mark Rolland; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). ANTOS, Christopher James; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). SIMONDS, Andrew Richard; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). JONES, Brian Michael; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). CLARKE, Simon Peter; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). SUNDERLAND, Edgar Mark; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). ROBINS, David Benjamin; c/o Microsoft Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). BOSE, Miko Arnab Sakhya Singha; c/o Microsoft

[Continued on next page]

(54) Title: COLLABORATIVE AUTHORING

FIG. 12



(57) **Abstract:** A collaborative authoring application provides an authoring environment in which two or more users can edit a document concurrently. Each user edits a copy of the document, sends updates to a master copy of the document, and receives updates from the master copy of the document. The authoring environment generally inhibits the users from providing conflicting editing instructions to the master copy of the document. For example, each user can generate a content lock about one or more data units within the document. The authoring environment may synchronize content locks automatically and content only at the request of the user.



Corporation International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (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, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,

ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

COLLABORATIVE AUTHORING

BACKGROUND

[0001] Traditional collaborative editing tends to be performed serially. Users take turns accessing a document, editing the document, and storing their edits. The accessing user may place a lock on the file to inhibit other users from editing the document when the accessing user is editing the document. The iterative editing process can cause delays since each user may wait for a turn at editing the document. In addition, the iterative editing process may be difficult to manage. For example, each user may need to keep track of who is editing which portions of the document, which version of the document is the most recent, and when the user will have a turn.

[0002] In other types of traditional collaborative editing, each user can edit a different copy of a document. Subsequently, all of the edited copies may be merged into a single document. This large scale merge also may cause delays, lead to numerous editing conflicts, and/or be difficult to manage. For example, the user responsible for merging the documents may be required to track the relationship between the documents. The user also may be responsible for resolving conflicts among two or more of the edited copies.

[0003] It is with respect to these and other considerations that the present disclosure has been made.

SUMMARY

[0004] This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended as an aid in determining the scope of the claimed subject matter.

[0005] A collaborative authoring application provides an authoring environment in which two or more users can edit a document concurrently. Each user edits a copy of the document, periodically sends updates to a master copy of the document, and periodically receives updates from the master copy of the document. The authoring environment generally inhibits the users from providing conflicting editing instructions to the master copy of the document. In other embodiments, the authoring environment can inhibit editing conflicts if the users edit the document at different times.

[0006] According to aspects of the disclosure, each document being authored can be divided into one or more data units. Each user can generate a content lock about one or more of the data units. Generating a content lock about a data unit inhibits other users from editing the locked data unit. In one embodiment, content locks can grow and/or shrink automatically as the user edits the document.

[0007] According to other aspects of the disclosure, the authoring environment synchronizes both content and metadata among the user copies and the master copy of the document. In general, the authoring environment synchronizes metadata automatically and synchronizes content only at the request of one of the users. In one embodiment, the metadata includes content locks.

[0008] These and other features and advantages will be apparent from a reading of the following detailed description and a review of the associated drawings. It is to be understood that both the foregoing general description and the following detailed description are explanatory only and are not restrictive of aspects as claimed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1 is a schematic block diagram illustrating an example authoring system having features that are examples of inventive aspects of the disclosure;

[0010] FIG. 2 is a schematic block diagram illustrating the authoring system of FIG. 1 in which a document stored on first computing device can include content and metadata in accordance with the principles of the present disclosure;

[0011] FIG. 3 is a flowchart illustrating an exemplary synchronization process by which an authoring system can synchronize a copy of a document stored on a user computing device with the master copy of the document in accordance with the principles of the present disclosure;

[0012] FIG. 4 is a schematic block diagram of a document having five units of data, the second of which is locked, in accordance with the principles of the present disclosure;

[0013] FIG. 5 is a schematic block diagram of the document of FIG. 4 in which a content lock has been added to the third data unit in accordance with the principles of the present disclosure;

[0014] FIG. 6 is a schematic block diagram of the document of FIG. 4 in which the content locks on the second and third data units have been released and a new content lock has been added to the fifth data unit in accordance with the principles of the present disclosure;

[0015] FIG. 7 is a schematic block diagram of a lock table in accordance with the principles of the present disclosure;

[0016] FIG. 8 is a flowchart illustrating an example synchronization processes implemented by an authoring application to share updates made by a user of the authoring application with other users in accordance with the principles of the present disclosure;

[0017] FIG. 9 is a flowchart illustrating another example synchronization process implemented by an authoring application to instantiate updates made by other users into the document being edited with the authoring application in accordance with the principles of the present disclosure;

[0018] FIG. 10 is a schematic block diagram of an authoring environment including a first computing device on which a master copy of a document to be authored is to be stored in accordance with the principles of the present disclosure;

[0019] FIG. 11 is a schematic block diagram of a user computing system configured to implement an authoring environment in accordance with the principles of the present disclosure;

[0020] FIG. 12 is a schematic block diagram of an authoring system illustrating a synchronization cycle implemented by the authoring application on the user computing device in accordance with the principles of the present disclosure;

[0021] FIG. 13 is a flowchart illustrating an operational flow for another exemplary synchronization process by which a user computing system can send and receive updates while editing a document in accordance with the principles of the present disclosure;

[0022] FIGS. 14-26 are schematic block diagrams of the caches stored on a first computing device, a first user computing device, and a second user computing device at different points in time during a collaborative authoring session in which a first user and a second user of the first and second user computing devices edit a document concurrently in accordance with the principles of the present disclosure; and

[0023] FIGS. 27-31 illustrate changes to an exemplary user interface of a first authoring application as a first user and a second user collaboratively author a document in accordance with the principles of the present disclosure.

DETAILED DESCRIPTION

[0024] In the following detailed description, references are made to the accompanying drawings that form a part hereof, and in which are shown by way of illustrations specific embodiments or examples. While the disclosure will be described in the general context of program modules that execute in conjunction with an application program that runs on an operating system on a computer system, those skilled in the art will recognize that the disclosure also may be implemented in combination with other program modules. The embodiments described herein may be combined and other embodiments may be utilized without departing from the spirit or scope of the present disclosure. The following detailed description is therefore not to be taken in a limiting sense, and the scope of the invention is defined by the appended claims and their equivalents.

[0025] Embodiments of the present disclosure provide an environment in which multiple users can collaboratively author a document. FIG. 1 illustrates an example authoring system 100 having features that are examples of inventive aspects of the disclosure. The authoring system 100 includes a storage device 120 storing a master copy of a document 150. In one embodiment, the storage device 120 can include a computing device. In another embodiment, the storage device 120 can include one or more storage devices (e.g., a network of computing devices).

[0026] The authoring system 100 also includes at least one user computing device 110 that is communicatively coupled to the storage device 120. Each of the user computing devices 110 can edit the document 150 by creating a user copy 155 of the document 150 and editing the user copy 155. The user copies 155 of the document 150 are synchronized when the user computing devices 110 periodically send to the storage device 120 updates to be shared with the other user computing devices and periodically obtain from the storage device 120 updates from other user computing devices.

[0027] As the term is used herein, a user computing device 110 includes any computing device that obtains a user copy of a document to be authored from a master copy of the document. The user computing device 110 can be different from the storage device 120 or can include a different user account implemented on the storage device 120. In one embodiment, a computing device that acts as a storage device 120 for one document may act as a user computing device 110 for a different document and vice versa.

[0028] In the example shown, four user computing devices 110A, 110B, 110C, and 110D are communicatively coupled to the storage device 120. In other embodiments, however, any number of computing devices 110 may be coupled to the storage device 120.

In the example shown, each user computing device 110A, 110B, 110C, 110D can send to the storage device 120 updates generated by the user of the user computing device and can request from the storage device 120 updates generated by the users of the other user computing devices. In one embodiment, the storage device 120 can be a server computing device and the user computing devices 110A, 110B, 110C, 110D can be client computing devices.

[0029] As shown in FIG. 2, the document 150 stored on the storage device 120 can include content 152 and metadata 154. Authoring applications 130 on the user computing devices 110 process and manipulate the content and metadata of the user copies 155 of the document 150. In some embodiments, metadata 154 can be stored separately from content 152. For example, content 152 can be stored in the document 150 and metadata 154 can be stored in a table (see FIG. 7) separate from the document 150. In other embodiments, however, the metadata 154 can be stored within the document 150.

[0030] In general, the user computing devices 110 can synchronize updates to the content 152 separately from updates to the metadata 154. In general, metadata updates 154 are automatically synchronized among the storage device 120 and user computing devices 110, whereas content updates 152 from each user computing device 110 are synchronized at the request of the respective user.

[0031] FIG. 3 is a flowchart illustrating an exemplary synchronization process 300 by which an authoring system 100 can synchronize a copy 155 of a document stored on a user computing device 110 with the master copy 150 of the document stored on the storage device 120. The synchronization process 300 initializes and begins at a start module 302 and proceeds to a first update operation 304.

[0032] The first update operation 304 synchronizes the metadata of the copy 155 with the metadata of the document 150 without requiring any interaction with the user. For example, the first update operation 304 can provide an exchange of metadata between the storage device 120 and the user computing device 110 at periodic time intervals. In one embodiment, the first update operation 304 provides for an exchange of metadata every few seconds. In other embodiments, however, the first update operation 304 can provide for the exchange of metadata to occur in shorter or longer time intervals.

[0033] A request module 306 determines whether a request to synchronize content has been received. For example, the request module 306 can determine whether a user of one of the user computing devices 110 has requested to share content changes with other users. In another embodiment, the request module 306 can determine whether a user of one

of the user computing devices 110 has requested to view content changes made by other users.

[0034] If the request module 306 determines that no request to synchronize content has been made, then the synchronization process 300 cycles back to the first update operation 304 and begins again. If the request module 306, however, determines that a request to synchronize content has been received, then the synchronization process 300 proceeds to a second update operation 308, which provides an exchange of content between the storage device 120 and the user computing device 110. The synchronization process 300 completes and ends at a stop module 310.

[0035] As the term is used herein, content updates 152 refer to any addition, deletion, and/or revision made to the substantive content of a document. For example, content updates for a word processing document can include added paragraphs (i.e., or sections thereof), deleted paragraphs (i.e., or section thereof), and/or revised paragraphs (i.e., or sections thereof). In another embodiment, content updates for a presentation document can include added, deleted, and/or revised pictures, text, animations, sounds, and other such data objects.

[0036] As the term is used herein, metadata updates 154 refer to any addition, deletion, and/or revision made to metadata of the document. Non-limiting examples of document metadata include content locks, presence information, and other such data. As will be discussed herein, content locks inhibit editing of content within the lock by users that do not own the lock. Presence information indicates which users have indicated an intention to edit the document as will be discussed in greater detail herein.

[0037] Referring to FIGS. 4-9, content locks inhibit editing conflicts by indicating which portions of a document have been claimed by another user. In some embodiments, the content locks can prevent (i.e., bar) a user from editing a portion of a document that has been claimed by another user. In other embodiments, however, the user can choose to break the content lock and edit the portion of the document. In such cases, the lock can warn the user that conflicts may arise when editing the locked section.

[0038] For example, FIG. 4 is a schematic block diagram of a document 400 having five units of data 410, 420, 430, 440, 450. In one embodiment, the document 400 is a word processing document and the units of data 410-450 are paragraphs of text. In another embodiment, the document 400 is a presentation document and the first unit of data 410 is a title or subject heading, the second unit of data 420 is a picture or other data object, and the remaining units of data 430, 440, 450 are blocks of text.

[0039] A first user can generate a lock about one or more of the units of data 410-450 of the document 400. In the example shown in FIG. 4, the second unit of data 420 has been locked by the first user, as indicated by the cross-hatching. By generating the content lock, the first user has indicated the first user intends to edit the unit of data 420. For example, the user can generate the lock when the user is actively editing the unit of data 420. In another embodiment, the user can generate the lock and then edit the unit of data 420 at a later time. As noted above, the authoring application managing the document 400 may prevent a user other than the first user from editing the locked data unit 420.

[0040] In general, a content lock generated by a first user is displayed to all other users who access the document during the life of the lock. In one embodiment, the first user also can view the content lock. Locks can be displayed using different types of indicia. For example, in one embodiment, the background of a locked unit of data may be colored, shaded, or patterned (e.g., see data unit 420 of FIG. 4). In another embodiment, the content (e.g., text, picture, shape, or other data object) of the locked unit may be colored, shaded, or patterned. In yet another embodiment, a box, bracket, or symbol can be displayed adjacent the locked unit of data to indicate the lock.

[0041] A content lock may be generated around one or more units of data within a document. In some embodiments, the content lock can be generated implicitly without express instruction from the user. For example, an authoring application may generate a content lock about a data unit when a user positions a cursor within, highlights, or otherwise selects the data unit. In another embodiment, an authoring application may generate a content lock about a data unit when the user begins editing the data unit. In another embodiment, an authoring application may generate a content lock about a data unit when a user saves edits made to the data unit. In some embodiments, the authoring application may generate content locks around multiple, inter-dependent data units within the document even if only one of the data units is claimed by a user. In other embodiments, a user may explicitly define a content lock about one or more data units.

[0042] In some embodiments, a content lock can grow to include additional units of data within the document and/or shrink to release one or more units of data within the document. In one embodiment, the content lock can grow to include only contiguous units of data. For example, as shown in FIG. 5, the content lock (e.g., see cross-hatching) around the second data unit 420 can expand to include the third data unit 430 if the user selects or edits the third data unit 430. A content lock or portion thereof generally can be released at the discretion of the user who generated the content lock. In some embodiments, the user

who generated the content lock releases the content lock by selecting or editing another data unit.

[0043] In one embodiment, the user must indicate the user has finished editing the data unit before the authoring application will release the lock. For example, the authoring application may release the lock only after the user has provided instructions to store the document. In another embodiment, the authoring application may release a lock after the user selects or begins editing a remote data unit, such as the fifth data unit 450 of FIG. 6. In another embodiment, the authoring application only releases a lock after both movement to a non-contiguous data unit and a finalization act (e.g., instructions to store the document). In other embodiments, a content lock can be released explicitly by an administrator or another user.

[0044] Referring to FIG. 7, lock metadata can be stored in a variety of different formats. For example, the lock metadata of FIG. 7 is stored in a table format 500. The lock table 500 of FIG. 7 includes a list of users, each of whom is identified with a user identifier (e.g., an identification number) that is uniquely assigned to the user. Data units to be locked are identified with unit identifiers (e.g., identification numbers) that are uniquely assigned to each data unit within a document. The lock table 500 associates the unit identifiers of the one or more data units to be locked with the user identifiers of the users who own the locks.

[0045] For example, in the lock table 500, data units 512 and 514 are associated with a first user 510. Other users, therefore, are inhibited from editing data units 512 and 514. Data unit 522 is associated with user 520. Other users, including the first user 510, therefore, are inhibited from editing data unit 522. The fourth user 540 has not locked any portion of the document and so is not associated with any unit identifiers. In other embodiments, however, lock metadata can be stored in a different format or within the document. For example, the lock table 500 can be arranged by unit identifier instead of by user identifier.

[0046] Presence metadata also can be stored in a variety of formats. For example, presence metadata can be stored in the lock table 500 of FIG. 7. In another embodiment, however, presence metadata can be stored in a separate table or in a different format. Presence metadata includes the user identifier of each user that is currently accessing the document or that has staked a claim (e.g., generated a content lock) on a data unit of the document. For example, a metadata table, such as the lock table 500, can store the user

identifier of each user having a claim to at least one data unit of the document. Like lock metadata, presence metadata can be synchronized automatically.

[0047] FIGS. 8 and 9 are flowcharts illustrating example synchronization processes 600, 700 by which content and locks can be synchronized among multiple computing devices. The first synchronization process 600 is implemented by an authoring application to share with other users updates made by a first user with the authoring application. The second synchronization process 700 is implemented by an authoring application to instantiate updates made by other users into the document being edited with the authoring application.

[0048] The first synchronization process 600 of FIG. 8 initializes and begins at a start module 602 and proceeds to a receive operation 604. The receive operation 604 receives an instruction from the first user to perform an editing operation on one or more data units within a document. Non-limiting examples of editing operations include adding, editing, and/or deleting data units (e.g., text, drawings, pictures, shapes, headings, etc.) or portions thereof, adding and/or deleting formatting of the data units, rearranging a sequence of the data units, and other such operations.

[0049] A first determination module 606 checks if the data unit on which the editing operation is to be performed is locked. For example, the first determination module 606 can access a lock table, such as the lock table 500 of FIG. 7, to determine whether a unit identifier assigned to the data unit is associated with any user. In another embodiment, the first determination module 606 can otherwise determine whether the data unit is locked.

[0050] If the first determination module 606 determines the data unit is not locked, then a generate operation 608 creates a content lock on the data unit. For example, in one embodiment, the generate operation 608 can add an entry to the lock table 500 of FIG. 7 associating the unit identifier of the data unit with the user identifier of the first user. A transmit operation 610 sends lock metadata indicating the existence of the newly generated lock to the device storing the master copy of the document. For example, the transmit operation 610 can send the lock metadata from a user computing device, such as computing device 110A of FIG. 1, to a storage device, such as storage device 120. In one embodiment, the lock metadata includes instructions that describe incremental edits (e.g., delta states) to the lock metadata. An implement operation 614 performs the editing operation on the data unit.

[0051] Alternatively, if the first determination module 606 determines the data unit is locked, then a second determination module 612 obtains ownership information for the

lock. In one embodiment, if the second determination module 612 determines the lock is owned by another user, then the synchronization process 600 completes and ends at a stop module 620 without performing the editing operation. In another embodiment (not shown), the synchronization process 600 may enable the first user to choose to break the lock and then proceed to generate operation 608. However, if the second determination module 612 determines the lock is owned by the first user, then the synchronization process 600 proceeds to the implement operation 614 discussed above. In another embodiment (not shown), the synchronization process 600 can determine whether the lock on any data units should be released.

[0052] A third determination operation 616 determines whether the first user has provided an instruction (either explicitly or implicitly) to share content changes to the data unit resulting from implementing the editing operation. For example, the third determination operation 616 may determine whether the first user has provided an instruction to store the document. In another embodiment, the third determination operation 616 may determine whether the first user has selected a "share changes" button or other such indicia. In another embodiment, the third determination operation 616 may determine an auto-save feature has been implemented by the application. If the third determination operation 616 determines instructions to share have not been provided, then the synchronization process 600 cycles back to the receive operation 604 and begins again.

[0053] If the determination operation 616 determines instructions to share have been provided, however, then the synchronization process 600 proceeds to a transmit operation 618, which sends content updates to the device storing the master copy of the document. Content updates indicate changes made to the content of the document by the editing operation. For example, the transmit operation 618 can send the content updates from a user computing device 110A of FIG. 1 to a storage device 120. In one embodiment, content updates include a newly generated version of the document. In another embodiment, content updates indicate incremental edits (e.g., delta states) to the content between the current user copy of the document and the master copy of the document. The synchronization process 600 completes and ends at the stop module 620 as discussed above.

[0054] The second synchronization process 700 of FIG. 9 initializes and begins at a start module 702 and proceeds to a receive operation 704. The receive operation 704 receives an update based on changes made to the document by other users. The update can include changes to the content and/or metadata of the document. In one embodiment, the

receive operation 704 can poll the device storing the master copy of the document and request updates. In another embodiment, the device storing the master copy can push changes to the authoring application.

[0055] A first determination module 706 determines whether the update includes any changes to the metadata of the document. For example, the first determination module 706 can determine whether any other user has established a new lock or released a lock on a data unit. In another embodiment, the first determination module 706 can determine whether another user has accessed the document.

[0056] If the first determination module 706 determines a metadata update has been received, then an implement operation 708 automatically instantiates the received metadata updates. For example, the implement operation 708 can cause the authoring application to update a lock table, such as lock table 500 of FIG. 7, based on the metadata update. However, if the first determination module 706 determines no metadata update has been received, then the second synchronization process 700 proceeds to a second determination module 710.

[0057] The second determination module 710 determines whether the update includes any changes to the content of the document. For example, the second determination module 710 can determine whether any data units have been added, deleted, revised, or moved. If the second determination module 710 determines a content update has not been received, then the second synchronization process 700 cycles back to the receive operation 704 to begin again. However, if the second determination module 710 determines a content update has been received, then the second synchronization process 700 proceeds to an indicate operation 712.

[0058] The indicate operation 712 alerts the first user that a content update is available for viewing. For example, the indicate operation 712 can display a graphic, text box, numerical indicia, or other message indicating the availability of the content update. In one embodiment, the indicate operation 712 provides information to the user indicating how many updates are available. Other information about the updates (e.g., user who made the updates, timestamp, etc.) also may be indicated.

[0059] A third determination module 714 determines whether the authoring application has received instructions from the first user (either explicitly or implicitly) to view and/or instantiate the content updates. For example, the third determination module 714 can determine whether the user has selected a view option (e.g., via a button, menu, or other interface tool) of the authoring application. In another embodiment, the third

determination module 714 may determine the user has selected an auto-instantiate option. If the third determination module 714 determines the authoring application has not received such instructions from the first user, then the second synchronization process 700 cycles back to the receive operation 704 and begins again.

[0060] If the third determination module 714 determines the authoring application has received instructions from the first user, then an instantiate operation 716 displays the content updates to the first user. For example, in one embodiment, the instantiate operation 716 can merge the available content updates into the first user's copy of the document. In another embodiment, the instantiate operation 716 can annotate the first user's copy of the document to indicate which content has changed. Additional annotations can indicate which users made which changes. In one embodiment, the instantiate operation 716 performs conflict resolution between the user copy and the updates. The second synchronization process 700 completes and ends at a stop module 718.

[0061] FIGS. 10-12 provide greater detail in how synchronization between the user copy and the master copy of the document is implemented by a user computing device. FIG. 10 is a schematic block diagram of an authoring system 800 including a storage device 820 on which a master copy of a document to be authored is to be stored. The authoring system 800 also includes at least one user computing device 810 communicatively coupled to the storage device 820.

[0062] The user computing device 810 includes an authoring application 812 configured to provide an authoring environment in which a user can create and/or manipulate a document to be authored. The user computing device 810 also includes a cache 814, a layer object ("LO") 816, and a synchronization manager ("sync manager") 818. The cache 814 stores a user copy of the document to be authored. The cache 814 also stores the metadata, including lock and presence metadata, associated with the document. Updates to the content and metadata of the document also can be stored in the cache 814.

[0063] The layer object 816 provides an interface between the authoring application 812 and the cache 814. The layer object 816 also provides an interface between the authoring application 812 and the sync manager 818. The sync manager 818 communicates with the storage device 820 and provides an interface between the storage device 820 and the cache 814. For example, the sync manager 818 can send updates to and obtain updates from the storage device 820 and the cache 814.

[0064] In general, an authoring environment having features that are examples of inventive aspects in accordance with the principles of the disclosure can be implemented on

a user computing device (e.g., a personal computer, a server computer, a notebook computer, a PDA, a Smartphone, or any other such computing device). A non-limiting embodiment of a user computing system 900 configured to implement an authoring environment is described herein with reference to FIG. 11.

[0065] In FIG. 11, the exemplary computing system 900 for implementing the principles of the disclosure includes a user computing device, such as user computing device 910. In a basic configuration, the user computing device 910 typically includes at least one processing unit 915 for executing applications and programs stored in system memory 920. Depending on the exact configuration and type of computing device 910, the system memory 920 may include, but is not limited to, RAM, ROM, EEPROM, flash memory, CD-ROM, digital versatile disks (DVD) or other optical storage devices, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or other memory technology.

[0066] System memory 920 typically stores an operating system 922, such as the WINDOWS® operating systems from MICROSOFT CORPORATION of Redmond, WA, suitable for controlling the operation of the computing device 910. System memory 920 also may include a document cache 926 in which a user copy 927 of a document can be stored. Metadata 929 of the document also can be stored within the user cache 926.

[0067] The system memory 920 also may store one or more software applications, such as authoring applications 924 for creating and editing documents. One non-limiting example of an authoring application 924 suitable for authoring documents in accordance with the principles of the present disclosure is MICROSOFT® OFFICE WORD authoring software from MICROSOFT CORPORATION of Redmond, WA. Other non-limiting examples of authoring applications include POWERPOINT® presentation software and VISIO® drawing and diagramming software, both also from MICROSOFT CORPORATION of Redmond, WA.

[0068] Computing device 910 also may have input device(s) 930, such as a keyboard, mouse, pen, voice input device, touch input device, etc., for entering and manipulating data. Output device(s) 935, such as a display screen, speakers, printer, etc., also may be included. These output devices 935 are well known in the art and need not be discussed at length herein.

[0069] The computing device 910 also may contain communication connections 940 that allow the device 910 to communicate with other computing devices, for example, the storage device 820 of FIG. 10, over a network in a distributed computing environment

(e.g., an intranet or the Internet). By way of example, and not limitation, communication device media 940 includes wired media such as a wired network or direct-wired connection, and wireless media, such as acoustic, RF, infrared and other wireless media.

[0070] FIG. 12 is a schematic block diagram of an authoring system 1000 illustrating the synchronization cycle implemented by the authoring application on the user computing device. The authoring system 1000 includes a storage device 1020 on which the document to be authored is to be stored and a user computing device 1010. The user computing device 1010 includes an authoring application 1012, a cache 1014, a layer object 1016, and a sync manager 1018. The cache 1014 of the user computing device 1010 is configured to store a base copy 1013 of the document to be authored, a working copy 1015, an upload copy 1017, and a download copy 1019. The cache 1014 also is configured to store metadata 1011 of the document.

[0071] FIG. 13 is a flowchart illustrating an operational flow for an exemplary synchronization process 1100 by which a user computing system, such as user computing device 1010, can send and receive updates while authoring a document. The synchronization process 1100 initializes and begins at a start module 1102 and proceeds to an obtain operation 1104. The obtain operation 1104 causes the sync manager 1018 of the user computing device 1010 to obtain from the storage device 1020 a download copy 1017 of the document to be authored. The download copy 1017 of the document reflects the state of the master copy of the document at the time the document is obtained from the storage device 1020. For example, the obtain operation 1104 can cause the sync manager 1018 to pull from the storage device 1020 a copy of the most recent version of the master copy of the document.

[0072] A generate operation 1106 creates a base copy 1013 and a working copy 1015 of the document based on the download copy 1017. The working copy 1015 of the document can be manipulated by the authoring application 1012 to allow the user to edit the document. Accordingly, the working copy 1015 reflects the current state of the document as edited by the user of the authoring application 1012. The base copy 1013 reflects the state of the document at the time just before the user begins editing the document. In one embodiment, the layer object 1016 generates the base copy 1013 of the document and the working copy 1015 and provides the working copy 1015 to the authoring application 1012.

[0073] To obtain content updates generated by other users, a poll operation 1108 causes the sync manager 1018 to contact the storage device 1020 periodically to obtain an

updated download copy 1017 of the document being authored. The download copy 1017 reflects the state of the master copy of the document at the time of download. The master copy is stored on storage device 1020 and may include changes made by other users since the base copy 1013 was generated. The sync manager 1018 stores any updated download copy 1017 of the document in the cache 1014.

[0074] An edit operation 1110 receives any instructions provided by the user to the authoring application 1012 to perform an editing operation on the document. The application 1012 makes the edits to the working copy of the document 1015. The authoring application 1012 also makes metadata changes as appropriate when implementing the editing operations.

[0075] At a first determination module 1112, the layer object 1016 checks the cache 1014 to determine whether any updates have been received during the poll operation 1108 (i.e., whether an updated download copy 1017 is available in the cache 1014). In one embodiment, the first determination module 1112 is performed at periodic intervals regardless of whether the user has edited the working copy of the document. In different embodiments, the first determination module 1112 can be performed, *e.g.*, every few seconds, milliseconds, or minutes.

[0076] If the layer object 1016 determines a new download copy 1017 is not available, then the synchronization process 1100 proceeds to a second determination module 1116, which will be discussed in greater detail below. If the layer object 1016 determines a new download copy 1017 is available, however, then the synchronization process 1100 proceeds to a sync operation 1114. In one embodiment, the sync operation 1114 automatically instantiates any metadata updates found in the download copy 1017 without requiring any interaction with the user of the authoring application 1012. For example, the layer object 1016 may automatically instantiates any lock metadata reflected in the download copy 1017 into the working copy 1015 of the document.

[0077] In some embodiments, the sync operation 1114 does not automatically instantiate any content updates from the download copy 1017 into the working copy 1015. Rather, the sync operation 1114 only instantiates content updates when the user of the authoring application 1012 provides instructions to integrate the content updates. In other embodiments, however, the sync operation 1114 automatically instantiates the content updates. In one embodiment, the sync operation 1114 causes the layer object 1016 to merge the updated download copy 1017 with the working copy 1015 when instructions to integrate are provided. In one embodiment, the sync operation 1114 is the same as the

second synchronization process 700 shown in FIG. 9. In other embodiments, however, other synchronization processes can be used to integrate the updates with the working copy 1015 of the document. The synchronization process 1100 proceeds from the sync operation 1114 to the second determination module 1116.

[0078] The second determination module 1116 determines whether the user made any changes to the working copy 1015 of the document or to the metadata 1011 of the document during the edit operation 1110. If the second determination module 1116 determines that no edits (i.e., neither content nor metadata) have been made, then the synchronization process 1100 cycles back to the poll operation 1108 and begins again. If the second determination module 1116 determines the user has made content edits (e.g., revising a data unit of the document) and/or metadata edits (e.g., locking a data unit of the document), however, then a share operation 1118 synchronizes the edits with the storage device 1020.

[0079] The share operation 1118 causes the sync manager 1018 of the user computing device 1010 to send automatically to the storage device 1020 any changes to the metadata of the document without requiring any interaction from the user of the authoring application 1012. In one embodiment, the share operation 1118 does not cause the sync manager 1018 to send content updates to the storage device 1020 automatically. Rather, the sync manager 1018 only sends content updates to the storage device 1020 when instructed to do so by the layer object 1016. Automatically pushing only metadata may enhance the user experience by increasing update efficiency, reduce network bandwidth usage, and provide conflict inhibition without requiring auto-syncing of content. In other embodiments, however, the share operation 1118 automatically pushes content as well as metadata. In one embodiment, the share operation 1118 is the same as the first synchronization process 600 of FIG. 8. In other embodiments, however, other synchronization processes can be utilized.

[0080] For example, when the user wishes to share content revisions made by the user, the user will instruct the authoring application 1012 to submit the revisions to the storage device 1020 for integration with the master copy of the document. In one embodiment, the authoring application 1012 will instruct the layer object 1016 to generate an upload copy 1019 of the document. Generally, the layer object 1016 generates the upload copy 1019 based on the working copy 1015. The sync manager 1018 attempts to push the upload copy 1019 to the storage device 1020. The synchronization process 1100 completes and ends at a stop module 1120.

[0081] In some embodiments, the share operation 1118 cannot send content updates to the storage device 1020 until all previous content updates received from the storage device 1020 have been instantiated into the working copy 1015 of the document. In such embodiments, the upload copy 1019 can be generated by merging the working copy 1015 of the document with the most recently obtained download copy 1017. In one embodiment, the sync manager 1018 polls the storage device 1020 for an updated download copy 1017 when the layer object 1016 attempts to generate an upload copy 1019.

[0082] If the sync manager 1018 fails to push the upload copy 1019 to the storage device 1020 (e.g., if the upload copy 1019 does not reflect the most recent changes instantiated into the master copy), then the sync manager 1018 retry with a revised update document 1019. For example, the sync manager 1018 may obtain a new download copy 1017 from the storage device 1020 and the layer object 1016 may generate a new upload copy from the working copy and the new download copy 1017. In one embodiment, the sync manager 1018 will continue to revise and push an upload copy 1019 to the storage device 1020 until successful.

[0083] Referring to FIGS. 14-31, the principles of the present disclosure can be better understood by walking through some example applications. FIGS. 14-26 represent changes made to a cache of a storage device, a cache 1214 a first user computing device, and a cache 1234 of a second user computing device during a collaborative authoring session in which a first user and a second user edit a document concurrently. In other embodiments, however, the first and second users can edit the document at different times using the same processes.

[0084] In FIG. 14, the first user has created a document and saved an initial draft of the document to a storage device as a master copy 1250. Alternatively, the first user may have accessed the master copy 1250 of an existing document. The second user has not yet accessed the document.

[0085] The master copy 1250 of the document has a corresponding master lock table. For the sake of convenience, in the example shown in FIGS. 14-26, the master lock table is displayed as part of the master copy 1250. In other embodiments, however, the master lock table can be stored separately from the master copy 1250 of the document. Similarly, for ease in understanding, locks associated with each copy of the document (e.g., base copy, working copy, download copy, and upload copy) stored on a user cache 1214, 1234 are shown in FIGS. 14-26 as part of the copy. In other embodiments, however, metadata may be separately stored for each document copy or for each user.

[0086] The master copy 1250 of the document includes a first data unit containing the text "Hello World" and a lock E1 on the first data unit. The lock E1 is owned by (i.e., assigned to) the first user. The content and metadata of the master copy 1250 are reflected in the base copy and download copy stored in the cache 1214 of the first user. Because the first user sent the initial draft of the document to the storage device, the upload copy stored in the cache 1214 also reflects the content and metadata of the master copy 1250.

[0087] Since sending the initial draft to the storage device, the first user has edited the working copy of the document to include a second data unit containing the text "Testing." A second lock E2 has been added about the second data unit. In one embodiment, the second lock E2 is a new lock assigned to the user. In another embodiment, the second lock E2 is an extension of the first lock E1. For example, a new unit identifier can be assigned to the first user in a metadata table stored in the cache 1214 of the first user. In one embodiment, the first user creates the second lock E2 by moving a cursor to start a new data unit (e.g., paragraph) and typing "Testing." The first lock E1 is not yet released because the first user has not yet indicated editing on the first data unit is final (e.g., by saving after moving the cursor).

[0088] In FIG. 15, the second user accesses the master copy 1250 of the document stored on the storage device. The second user obtains a download copy of the document and stores the download copy into the cache 1234 of the second user. The download copy includes the content and metadata of the master copy 1250. The download copy does not include the most recent content and metadata contained in the working copy of the first user. The layer object of the second user generates a base copy and a working copy of the document based on the download copy. The authoring application utilized by the second user displays the content and any locks of the working copy to the second user.

[0089] In some embodiments, the second user can edit the document even before a coauthoring session has been fully established (i.e., before the first and second user computing devices learn of the presence of the other). For example, the second user is allowed to edit the second working copy of the document while the first and second user computing devices exchange greetings (i.e., handshake). Typically, metadata is not exchanged automatically before the first and second computing devices learn of the existence of one another.

[0090] In FIG. 16, the first and second computing devices each have discovered the existence of the other. Accordingly, each of the first and second computing devices has begun transmitting metadata updates automatically to each other. As shown in FIG. 16, the

first authoring application provides updated lock metadata (e.g., the second lock) from the lock table of the first user to the storage device 1220 (e.g., via a layer object and sync manager). The authoring application does not provide updated content information to the storage device 1220, however, because no instructions to share content updates have been received from the first user.

[0091] In FIG. 17, the storage device has integrated the updated lock metadata (e.g., the second lock E2) into the master metadata table of the master copy 1250. The lock metadata has not yet been obtained by the second user computing device. In FIG. 18, the second user computing device has polled the storage device and obtained a new download copy of the master copy 1250. A new download copy also has been generated (e.g., via a layer object) on the first cache 1214 based on the upload copy. In another embodiment, the first user computing device also may have polled the storage device and obtained a new download copy of the master copy 1250. Both of the new download copies reflect the updated lock information (e.g., lock E2) from the working copy of the first user.

[0092] In FIG. 19, the second computing device generates (e.g., via a layer object) a new working copy based on the new download copy. Accordingly, the new working copy of the second computing device contains the second lock E2 on the second data unit. The new working copy does not include the second data unit itself, however. In one embodiment, the working copy of the second user computing device displays a blank/empty data unit and indicates that the blank/empty data unit is locked. In another embodiment, the working copy does not display the second data unit at all and the lock E2 is not indicated to the second user until the second data unit is visible to the second user.

[0093] In FIG. 20, the second user adds a third data unit to the document and a third lock E3 around the third data unit (e.g., in the memory of the authoring application). The third data unit contains the text "Goodbye World". The second authoring application generates (e.g., via a layer object) a working copy of the document including the third data unit and the third lock E3. The second authoring application also automatically provides (e.g., via a sync manager) the third lock E3 to the storage device. However, the second authoring application does not provide the third data unit to the storage device.

[0094] Also in FIG. 20, the first user releases the first lock E1 after the first user indicates the first user has finished editing the first data unit. For example, the first user can provide instructions to store the document when the cursor of the first user is positioned within the second data unit. In one embodiment, the second data unit may be located remote from the first data unit. In the example shown, the finalization indication

provided by the first user does not indicate the first user wishes to share content changes made by the first user. Accordingly, the first authoring application does not generate (e.g., via a layer object) an upload copy containing content changes for distribution to the storage device.

[0095] In FIG. 21, the third lock E3 has been stored in the master lock table of the master copy 1250 of the document on the storage device. The first user computing device subsequently polls (e.g., via a sync manager) the storage device to check for updates. The first user computing device obtained (e.g., via a sync manager) a new download copy of the document from the storage device based on the master copy 1250. The new download copy contains the third lock E3. The first authoring application automatically instantiates (e.g., via a layer object) the third lock E3 into the working copy of the first authoring application. As noted above, the third lock E3 optionally may be displayed to the first user using the user interface of the first authoring application.

[0096] In FIG. 22, the second user releases the third lock E3 when the second user indicates the second user has finished editing the third data unit. The third lock E3 is removed from the metadata table of the second user. In the example block diagram shown in FIG. 22, the third lock E3 is removed from the metadata tables corresponding to the working copy of the second authoring application.

[0097] In the example shown in FIG. 22, the finalization indication provided by the second user indicates the second user wishes to share content changes with other users. For example, the second user may remove a cursor from the third data unit and then provide instructions to store the document. Accordingly, the second authoring application may generate an upload copy containing the content changes based on the working copy. In one embodiment, the upload copy also can be based on any content updates (e.g., in this case none) found in the download copy that had not yet been instantiated into the working copy. In the example shown, the upload copy includes the third data unit including the phrase "Goodbye World."

[0098] In FIG. 23, the storage device obtains the third data unit from the upload copy of the second user computing device and merges the third data unit into the master copy 1250 of the storage device. The lock table associated with the master copy 1250 indicates the third lock E3 has been released. As shown in FIG. 24, upon a successful push to the storage device, the second authoring application copies (e.g., via a layer object) the upload copy to the download copy and the base copy of the second cache 1234.

[0099] In FIG. 25, the first user computing device obtains (e.g., via a sync manager) a new download copy of the master copy 1250 of the document. The first authoring application determines (e.g., via a layer object) the new download copy has been obtained and determines the new download copy includes content changes. The authoring application automatically merges any new locks (none in this case) into the working copy of the cache 1214. The user interface of the first authoring application may indicate content updates are available for viewing.

[0100] In FIG. 26, the first user chooses to instantiate the content updates available in the cache 1214 into the working copy of the document. Accordingly, the authoring application merges the download copy stored in the cache 1214 with the working copy of the first authoring application. The third data unit is added to the working copy of the first authoring application and is displayed to the first user. The third lock E3 is released in the metadata table associated with the working copy of the first authoring application. The first user, therefore, may instruct the first authoring application to edit the third data unit.

[0101] Referring now to FIGS. 27-31, the user experience during the coauthoring session is described in detail below. FIGS. 27-31 illustrate changes to an exemplary user interface 2000 of a first authoring application as a first user and a second user collaboratively author a document 2010. The first user, Jane Doe, edits the document 2010 using the first authoring application and the second user, John Doe, edits the document 2010 using a different authoring application.

[0102] The document 2010 includes a first data unit 2012, a second data unit 2014, a third data unit 2016, and a fourth data unit 2018. The first user places a content lock 2022 around the first data unit 2012. For example, in one embodiment, the first user may move her cursor 2015 to the first data unit 2012. In another embodiment, the first user may provide instructions to perform one or more editing operations on the first data unit 2012.

[0103] In FIG. 28, the first user has added text to the first data unit 2012. The lock on the first data unit inhibits the second user from editing the first data unit 2012. In one embodiment, the first lock 2022 prevents the second user from editing the first data unit 2012. In another embodiment, the first lock 2022 warns the second user against editing the first data unit 2012, but does not bar editing by the second user. In yet another embodiment, the first lock 2022 merely indicates the first user has expressed an intention to edit the first data unit 2012.

[0104] In FIG. 29, the second user has placed a content lock 2024 about the third data unit 2016. Because the second user has chosen to not share his changes, the second

lock may be visible on the user interface 2000 of the first authoring application, but any changes made to the third data unit 2016 by the second user are not visible yet. In the example shown, the second lock 2024 is indicated by the user interface 2000 by displaying a bracket indicating the scope of the second lock 2024. The second lock 2024 also can include a nametag indicating the owner of the second lock 2024. In FIG. 29, the nametag includes text identifying the second user and an icon. The icon can further identify a user by color, symbol, or other graphic.

[0105] In FIG. 30, the first authoring application receives content updates for the document 2010. The second user has chosen to share content changes with other users and accordingly merged content updates into a master copy of the document 2010 (e.g., via one of the synchronization processes described above). A sync manager of the first authoring application obtained the updates from the master copy of the document 2010 by polling the device storing the master copy.

[0106] The user interface 2000 of the first authoring application displays an alert 2030 to the first user indicating the availability of the content update. In the example shown, a text box 2030 appears in the user interface 2000. In another embodiment, the second lock or text within the second lock can change color, pattern, or formatting. In another embodiment, an update counter increments when each new update is available.

[0107] The first user can choose to view/instantiate the changes made by the second user or to ignore the content updates. The first user can continue to edit any unlocked data units within the document 2010 while ignoring the content updates. In one embodiment, the first user must instantiate the content updates before merging her updates with the master copy. In FIG. 31, the first user instructs the first authoring application to instantiate the content updates into the document 2010 displayed by the user interface 2000 of the first authoring application. Content changes made by the second user are now visible to the first user.

[0108] Embodiments of the disclosure may be implemented as a computer process (method), a computing system, or as an article of manufacture, such as a computer program product or computer readable media. The processes (programs) can be implemented in any number of ways, including the structures described in this document. One such way is by machine operations, of devices of the type described in this document. Another optional way is for one or more of the individual operations of the methods to be performed on a computing device in conjunction with one or more human operators performing some of

the operations. These human operators need not be collocated with each other, but each can be only with a machine that performs a portion of the program.

[0109] The computer program product may be a computer storage media readable by a computer system and encoding a computer program of instructions for executing a computer process. The computer program product may also be a propagated signal on a carrier readable by a computing system and encoding a computer program of instructions for executing a computer process. The term computer readable media as used herein includes both storage media and communication media.

[0110] Those skilled in the art will appreciate that the disclosure may be practiced with other computer system configurations, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers, and the like. The disclosure may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices. Generally, program modules include routines, programs, components, data structures, and other types of structures that perform particular tasks or implement particular abstract data types.

WE CLAIM:

1. A method of facilitating collaborative editing of a document (2010) between a first user at a first computing device (110, 810, 910, 1010) and at least one other user, the method comprising:

accessing a user copy (155) of the document at the first computing device (110, 810, 910, 1010) including accessing content (152) of the document and accessing metadata (154, 1011) associated with the document, the content (152) of the document being organized into units of data (410-450);

receiving metadata updates periodically from a storage device (120, 820, 1020), the metadata updates including any locks on the units of data (410-450) of the document that are associated with the other user;

instantiating the received metadata updates automatically including storing the locks on the units of data (410-450) of the document with the metadata (154, 1011) associated with the document, wherein storing the locks with the metadata associated with the document prevents the first user from editing any units of data that have locks associated with the other user; and

receiving content updates periodically from the storage device (120, 820, 1020), the content updates including any modifications made to the content (152) of the document by the other user.

2. The method of claim 1, wherein accessing the user copy of the document at the first computing device comprises:

obtaining from the storage device a download copy of the document indicating a master state of the document;

generating a base copy and the user copy of the document based on the download copy;

downloading from the storage device the metadata associated with the document; and

storing the metadata on the first computing device.

3. The method of claim 1, wherein receiving content updates periodically from the user computing device comprises retrieving a download copy of the document from the storage device, wherein the download copy indicates a master state of the document.

4. The method of claim 3, further comprising instantiating the received content updates only when instructed by the first user including merging the retrieved download copy of the document with the user copy of the document.
5. The method of claim 1, further comprising:
 - receiving at the first computing device an indication of editing of one of the units of data of the document by the first user;
 - updating the metadata stored on the first computing device to lock the unit of data being edited by the first user, wherein locking the unit of data associates the unit of data with the first user; and
 - transmitting to the storage device automatically a metadata update based on the updated metadata stored on the first computing device.
6. The method of claim 5, further comprising:
 - transmitting a content update to the storage device indicating modifications made to the unit of data edited by the first user only when instructions to transmit the content update are received from the first user.
7. The method of claim 6, wherein transmitting the content update to the storage device comprises:
 - receiving instructions from the first user to transmit the content update to the storage device;
 - instantiating any content updates received previously from the storage device into the user copy of the document;
 - generating an upload copy of the document based on the user copy of the document including any instantiated content updates; and
 - transmitting the upload copy of the document to the storage device.
8. The method of claim 7, wherein receiving instructions from the first user to transmit the content update to the storage device comprises receiving instructions from the first user to store the document on the first computing device.

9. The method of claim 1, wherein receiving metadata updates periodically from the storage device comprises receiving presence information indicating identities of other users that are selected from a group consisting of users accessing the document and users associated with locks stored in the metadata associated with the document.

10. An authoring environment implemented on a first computing device (110, 810, 910, 1010) to enable coauthoring of a document among users, the first computing device being configured to communicate with at least a storage device (120, 820, 1020), the collaborative authoring environment comprising:

- a. a cache (814, 926, 1014, 1214, 1234) configured to store a base copy (1013) of the document, a download copy (1019) of the document, an upload copy (1017) of the document, a working copy (1015) of the document, and a lock table (500);
- b. an authoring application (812, 924, 1012) configured to manipulate the working copy (1015) of the document including manipulating content (152) of the working copy and manipulating metadata (154, 1011) of the working copy (1015);
- c. a layer object (816, 1016) that provides an interface between the authoring application (812, 924, 1012) and the cache (814, 926, 1014, 1214, 1234), the layer object (816, 1026) being configured to update the lock table (500) stored within the cache (814, 926, 1014, 1214, 1234) automatically with updated metadata when the authoring application manipulates the metadata of the working copy (1015), the layer object also being configured to generate the upload copy (1017) of the document based on the working copy of the document when instructions to generate the upload copy are received from the authoring application, wherein the upload copy includes content updates obtained from the download copy (1019) of the document; and
- d. a sync manager (818, 1018) configured to manage interactions between the cache (814, 926, 1014, 1214, 1234) and the storage device (120, 820, 1020), the sync manager being configured to push the updated metadata from the lock table to the storage device automatically and being configured to push the upload copy (1017) of the document from the cache to the storage device when instructions to provide the updates to the content are received from the authoring application.

11. The authoring environment of claim 10, wherein the sync manager also is configured to obtain any updated download copy periodically from the storage device and to store the updated download copy within the cache of the first computing device; wherein the updated download copy includes content updates generated by other users.
12. The authoring environment of claim 11, wherein the authoring application provides alerts indicating an updated download copy is available for viewing at the first computing device.
13. The authoring environment of claim 11, wherein the authoring application is configured to receive instructions from a user of the authoring application to view any content updates contained in the updated download copy.
14. The authoring environment of claim 10, wherein the authoring application includes a word processing application.
15. A computer readable storage medium storing computer executable instructions, which perform a method of authoring a document when executed by a computing device, the method comprising:
 - receiving at a first authoring application (812, 924, 1012) instructions to manipulate a first unit of data (410-450) of a document being edited by the first authoring application, the first authoring application being utilized by a first user and being implemented on a first computing device (110, 810, 910, 1010);
 - generating a content lock on the first unit of data (410-450) automatically after receiving the instructions to manipulate by associating the first unit of data with the first user in a lock table (500), the lock table being accessible to a plurality of other authoring applications, wherein only the first user can manipulate the first unit of data when the first unit of data has the content lock;
 - manipulating the first unit of data (410-450) of the document with the first authoring application in accordance with the instructions to manipulate; and
 - releasing the content lock on the first unit of data (410-450) only when instructions to release the content lock are provided to the first authoring application by the first user, wherein releasing the content lock includes disassociating the first unit of data from the first user in the lock table.

16. The computer readable storage medium of claim 15, wherein the method further comprises:

expanding the content lock to encompass additional units of data when instructions from the first user to manipulate the additional units of data are received by the first authoring application, wherein expanding the content lock includes associating the additional units of data with the first user in the lock table.

17. The computer readable storage medium of claim 15, wherein releasing the lock on the first unit of data comprises releasing the lock on the first unit of data when the first user provides to the first authoring application an indication that the first user has finalized edits to the first unit of data.

18. The computer readable storage medium of claim 15, wherein the lock table is stored on a remote storage device accessible by the other authoring applications.


19. The computer readable storage medium of claim 16, wherein the method further comprises:

transmitting automatically any lock updates to the lock table stored on the storage device.

20. The computer readable storage medium of claim 15, further comprising:

retrieving updates to the lock table from the storage device, the updates to the lock table having been made by at least one of the other users;

instantiating the updates to the lock table into the document automatically.






FIG. 3






FIG. 4




(T1)

FIG. 5



(T2)

FIG. 6



(T3)

FIG. 7




FIG. 8




FIG. 9





FIG. 10






FIG. 12




FIG. 13









FIG. 14



































FIG. 27




FIG. 28




FIG. 29




FIG. 30





FIG. 31



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2008/081456**A. CLASSIFICATION OF SUBJECT MATTER****G06F 17/21(2006.01)i, G06F 17/00(2006.01)i, G06F 17/24(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC : G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean Utility models and applications for Utility models since 1975
Japanese Utility models and applications for Utility models since 1975Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKIPASS(KIPO internal) "collaborative, authoring, metadata, update, lock, document, periodically, copy"**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 06983416 B1 (JIN-HAK BAE et al.) 3 JANUARY 2006 see abstract, column 5 line 6 - 41, column 6 line 57 - column 7 line 35, Figures 6, 9A-10.	1-20
A	US 2007/0186157 A1 (RICHARD P. WALKER et al.) 9 AUGUST 2007 see abstract, paragraph[0371] - paragraph[0392], Figures 16, 18-19.	1-20
A	US 2004/0107224 A1 (RAJENDRA KUMAR BERA) 3 JUNE 2004 see abstract, paragraph[0017] - paragraph[0025], Figures 1-2.	1-20

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

* Special categories of cited documents:
 "A" document defining the general state of the art which is not considered to be of particular relevance
 "E" earlier application or patent but published on or after the international filing date
 "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)
 "O" document referring to an oral disclosure, use, exhibition or other means
 "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
 "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
 "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
 "&" document member of the same patent family

Date of the actual completion of the international search

31 MARCH 2009 (31.03.2009)

Date of mailing of the international search report

31 MARCH 2009 (31.03.2009)

Name and mailing address of the ISA/KR


 Korean Intellectual Property Office
 Government Complex-Daejeon, 139 Seonsa-ro, Seo-gu, Daejeon 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

AN, BYUNG IL

Telephone No. 82-42-481-8471



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2008/081456

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6983416 B1	03.01.2006	KR 2001094087 A KR 10-0331685 B1	31.10.2001 09.04.2002
US 2007-186157 A1	09.08.2007	None	
US 2004-0107224 A1	03.06.2004	None	