



(51) International Patent Classification:

G06F 9/54 (2006.01)

G06Q 10/10 (2012.01)

G06F 8/33 (2018.01)

(21) International Application Number:

PCT/US2019/038823

(22) International Filing Date:

25 June 2019 (25.06.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/058,154

08 August 2018 (08.08.2018)

US

(71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: MOWATT, David; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Red-

mond, Washington 98052-6399 (US). JIMENEZ SALGADO, Rolando; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). CAMPBELL, Johnny S.; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). CHILAKAMARRI, Venkat Pradeep; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). BALZER, Andreas; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). WOULFE, Muiris; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). O'DRISCOLL, Stephen; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(74) Agent: MINHAS, Sandip S. et al.; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(54) Title: VISUALIZATION PLATFORM FOR REUSABLE DATA CHUNKS




FIGURE 1

(57) Abstract: Techniques for rendering shared data include receiving, from a data store by a computing device, data indicative of a persistent data object. The persistent data object is associated with a class indicative of a data type for information contained in the persistent data object. The persistent data object is operable to be inserted in a file generated by an application executing on the computing device. The visualization logic associated with the class may be received by the computing device. The information in the persistent data object is rendered using the selected method of visualization, and other information in the document is rendered using native rendering capabilities of the application.



(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).

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 (Art. 21(3))*

VISUALIZATION PLATFORM FOR REUSABLE DATA CHUNKS

BACKGROUND

[0001] Today's workforce accomplishes many tasks using a variety of devices, software applications, and web-based services. A user may, for example, use a phone to capture data and take notes, and share information with colleagues using one or more applications and services. The user may use the phone or another computing device to work on the information and develop other work products during the course of a project, using the same or different applications and services.

[0002] A number of documents may be generated and developed during the course of a project. Various versions of these documents may be generated. Finding documentation, past or present, about a topic can be time consuming and difficult if not impossible. Furthermore, some parts of documents may be copied or cut and pasted into other documents of the same or different applications. If the original source that was copied was incorrect or is later updated, then all downstream users who copied the original source will have incorrect or obsolete information.

[0003] Some systems provide the ability to insert or embed information into documents. However, such systems do not support sharing of the information or provide any way of updating the embedded information in a dynamic manner with control of the changes and in a way that the information can be readily shared. Furthermore, users typically use a variety of applications and devices, and some pieces of information must be modified or translated to be able to rendered or otherwise consumed by different applications and devices.

[0004] These shortcomings can be inconvenient for users and cause significant inefficiencies with respect to the utilization of computing resources. For example, opening and interacting with multiple applications to view and share relevant information regarding a particular context or topic can result in the inefficient utilization of processor cycles, memory, battery power, and network bandwidth. Similarly, the inability to accurately view shared information that was originated using other applications can lead to significant rework and delays due to manual translation and updating. Furthermore, the inability to view and update pieces of information across devices and applications without manual intervention can cause inefficient computer interactions that increase the frequency of errors which, in turn, cause further inefficiencies with respect to the utilization of computing resources. Given these drawbacks and others, there still exists a need for tools that can

efficiently render and update shared information in a dynamic manner.

[0005] It is with respect to these and other technical challenges that the disclosure made herein is presented.

SUMMARY

5 **[0006]** The present disclosure describes various embodiments that enable the rendering of a segment or chunk of information which can be shared and reused among users and documents. The rendering of the chunks of information allow the chunks to be more readily used and reused as a number of related or offshoot documents are generated and developed around the chunks. In various embodiments, the chunks of data may be allowed to be
10 rendered using basic text/table/picture forms of rendering and into more scenario-specific forms, while allowing the underlying chunks to be shared and used among multiple devices and applications.

[0007] In an embodiment, the visualization of one or more chunks from a chunk data store may be rendered in the context of a document file within the application using its native
15 rendering capabilities. In another embodiment, a visualization of one or more chunks may be rendered independent of the application.

[0008] In one embodiment, a visualization mechanism for viewing, editing, and otherwise accessing chunks may be provided. The visualization mechanism may be incorporated into an application or implemented as a standalone application or process. In some embodiments,
20 the chunks may include or may be linked to logic that enables an application to render the chunk information without the need for pre-existing functionality in the application or locally on the user device. In some embodiments, the logic may be configured to allow visualization of the chunk content. In one embodiment, the logic can provide a basic visualization for applications that are not configured to otherwise render the content of the
25 chunk. The logic may be associated with a chunk type. For example, a table type or a graphic type may each have associated logic that allows for rendering of the chunk type.

[0009] The associated chunk visualization logic may be implemented as code that can run on the user device or at the storage service when visualization is needed. For example, if the data chunk could be rendered as a table, and the local device is attempting to render the table
30 using Notepad which may not have a built-in or native capability to fully render a table, then Notepad may render the table using logic provided for the chunk. The logic may be co-located with the chunk at the storage service, or available for download to the user device. Otherwise, the Notepad application may continue to render the table using its limited rendering capabilities. For example, the contents of the table may be rendered as a list with

delimiters.

[0010] In some embodiments, the chunk rendering functionality may provide a plurality of visualization choices that are based on possible renderings based on the schema of the chunk, and optionally what is desirable based on the semantic match between the chunk and the visualization type. The visualization choices may also be based on what display options that the host application can support. The visualization choices may further be based on visualizations shared with the user, visualizations that the user acquired, visualizations that the group of the user acquired, visualizations that a parent group acquired, visualizations temporarily accessible through subscriptions, and other factors.

[0011] The techniques disclosed herein can improve user interaction with a computing device, which can increase a user's productivity and help reduce the number of inadvertent, obsolete, or otherwise incorrect inputs. Also, by providing more accurate and more contextually relevant information, a system can operate more efficiently with respect to the use of memory, processing resources, network resources, etc. Having more accurate and more contextually relevant data can reduce the number of emails and other forms of communication that may be used to clarify and correct inaccurate information. Other technical effects other than those mentioned herein can also be realized from implementations of the technologies disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] The Detailed Description is described with reference to the accompanying figures. The same reference numbers in different figures indicates similar or identical items. References made to individual items of a plurality of items can use a reference number with a letter of a sequence of letters to refer to each individual item. Generic references to the items may use the specific reference number without the sequence of letters.

[0013] FIGURE 1 illustrates a diagram illustrating the generation and use of chunks according to various embodiments.

[0014] FIGURE 2 illustrates an example table showing the contents of a chunk.

[0015] FIGURE 3 illustrates an example table showing the contents of a chunk.

[0016] FIGURE 4 illustrates an example table showing the contents of a chunk.

[0017] FIGURE 5 illustrates an example table showing the contents of a chunk.

[0018] FIGURE 6 illustrates a block diagram for identifying chunks and providing them to an application.

[0019] FIGURE 7 illustrates a block diagram for identifying chunks and providing them in an organization.

[0020] FIGURE 8 illustrates a block diagram for updating chunks.

[0021] FIGURE 9 illustrates a block diagram for updating chunks.

[0022] FIGURE 10 illustrates aspects of a routine for enabling aspects of the techniques disclosed herein as shown and described below.

5 [0023] FIGURE 11 illustrates aspects of a routine for enabling aspects of the techniques disclosed herein as shown and described below.

[0024] FIGURE 12 illustrates aspects of a routine for enabling aspects of the techniques disclosed herein as shown and described below.

[0025] FIGURE 13 is a computer architecture diagram illustrating an illustrative computer
10 hardware and software architecture for a computing system capable of implementing aspects of the techniques and technologies presented herein.

DETAILED DESCRIPTION

[0026] Described herein are systems and methods for generating and using chunks of data, memes, and information in formats and schemas that allow for the re-use or extension of
15 that information in a plurality of document types and applications. As used herein, a chunk of information, or a chunk, may generally be defined as a set of modifiable and reusable data that has associated metadata and may be uniquely identified. A chunk may further be composed of other chunks. The chunk of information may be referred to herein as a persistent data object. These terms may be used interchangeably. A chunk may include any
20 type of data such as text, graphic, table, or any other type of information that may be inserted into a document of any type, such as a word processing document, spreadsheet, presentation, and the like.

[0027] A chunk of information as disclosed herein may be configured to be created, stored, updated, sent and received, and inserted into documents as independent objects. When
25 inserted into documents, the chunk of information is incorporated into the document as other information may be inserted into the document such as in a cut and paste or copy and paste operation. However, the inserted chunk of information may continue to be independently identifiable and tracked. In other words, the chunk of information exists independently of the document in which it is incorporated. In one embodiment, when the chunk of
30 information is updated, either within the inserted document, or independently in another document that may be updated by an application that may be executing in the same or a different computer, the updates may be provided to all documents that have incorporated the chunk of information, regardless of the document type and application. By providing such a mechanism, a chunk of information may be consistently used without becoming obsolete

or otherwise contain stale or erroneous information. Additionally, because the chunk of information may be compatible with a plurality of document types and applications, chunks and their updates may be provided without regard to the specific type of document where they reside. Furthermore, the identity of a chunk of information may be tracked even when the contents have been updated.

[0028] As an example, a user may generate notes in a To-Do list application. The notes may then be evolved into survey questions in an Excel spreadsheet, a live survey run in an application on an iPhone, and finally into survey results in a PowerPoint presentation. A portion of the notes may be defined as a chunk that may in turn be incorporated into the Excel spreadsheet, survey, and the PowerPoint presentation. Various embodiments include mechanisms to track changes to chunks so that individual changes made to a particular chunk in one application may be applied wherever the chunk is incorporated. Furthermore, the changes may be made to the other documents based on the properties of the corresponding document types and their respective applications. In some embodiments, chunks may optionally include or be associated with functionality/logic for responding intelligently to and/or protecting edits. For example, a survey chunk can be configured to wipe results if a question is changed. The chunk may communicate with the displaying application to ensure that the application prompts the user as to whether to update the question and delete the results, or keep the results and discard the edit. Alternatively, the chunk may access or open the chunk management function to perform the user prompt.

[0029] The disclosed embodiments thus provide a way for a user to make updates in one document or application that can selectively be reflected where a chunk is used in other documents or applications. Additionally, in some embodiments the chunk may be associated with metadata such as data history and schema history. This information may be used to track the history of a chunk and its revisions. In some embodiments, the applications can each have schemas that they support, which can be declared or inferred from the context of the chunk.

[0030] For example, a To-Do list application which can support a single textual field per chunk may render “What is the Business opportunity size” as a note but may not properly render “What is the Business opportunity size SMALL | MEDIUM | LARGE” as a note, while a survey questions list in Excel that can support complex schema and multiple fields may display “What is the Business opportunity size” in one column and “SMALL | MEDIUM | LARGE” in the next column. A survey application may further recognize that the “SMALL | MEDIUM | LARGE” field is actually a multiple-response field and may

render a drop-down menu to allow selection by the user.

[0031] In various embodiments, when the data in the chunk is updated or augmented, for example by changing the question or by responding to the survey, versioned data may be added to the chunk. Each of the applications and documents that use the chunk may determine whether to incorporate the new data, and also whether the application can render this new data or an older version of the data, and if so, which aspects of the schema that they can support.

[0032] Additionally, chunks of information may be rendered in a plurality of ways that are independent of a particular rendering format, as long as the type of data can be rendered. In some embodiments, a visualization platform may be implemented that provides various rendering capabilities such as graphing and charting that can render the information in a chunk without requiring any particular application to natively render the information.

[0033] In an embodiment, a computing environment that implements chunks may be implemented using one or more components. In one embodiment, the components may execute on the same physical computing device. In other embodiments, the components may be distributed, for example in a client-server architecture.

[0034] FIGURE 1 illustrates a system 100 for enabling the generation, storage, and updating of chunks from multiple applications 141 and doing so while enabling a user to work within a contextual environment of each application 141. In this example, a user can interact with an individual application 141 to create and edit documents, and view and add or edit chunk content that is inserted in a particular type of file, e.g., a word processing document, a spreadsheet document, etc. The applications 141 may each be configured to display an authoring pane 191 and a chunk pane 192. The content of each file is displayed in the authoring pane 191. A user can select portions of content displayed in the authoring pane 191. The selected portions can be selected as chunks in a chunk pane 192 of any application to create and save new chunks. The chunk pane 192 may also be used to view available chunks for selection and insertion into the application. The chunk pane 192 of each application 141 may be synchronized to enable a user to work within each application 141 and view a current version of a document, e.g., a presentation, they are editing while also viewing chunk content specific to each application or available to each application.

[0035] The content in the chunk pane 192 can be used to generate chunk data 151 having content and metadata. In some configurations, the chunk data 151 can be in the form of a table, file, a multi-frame image file, a video file, or any other suitable format. In some embodiments, the chunk pane may display a tab or similar mechanism for each revision so

that the user can revert back to older or newer versions. The chunk pane may also provide a mechanism to easily see the differences between revisions.

[0036] In some embodiments, a system 100 can identify or generate chunk data for use by multiple users and applications. For example, a computing device 106A or a storage service
5 110 and management service 120 can identify chunk content and store the chunk content.

[0037] Chunk data 151 can include text, images, media or any other form of data. The chunk data 151 can include data that is stored within storage service 110 including a data store 136 and managed by management service 120 comprising a management module 138. In some embodiments, chunk data 152 can include visualization logic that is integrated into,
10 or linked to, the chunk data 151.

[0038] The chunk data 151 can be communicated to any number of computing devices 106, referred to herein as computing devices 106B-106N, from a first computing device 106A or the service 110 via a network 108. Each computing device 106B-106N associated with a recipient can display the chunk data 151 on a user interface 195 (195A-195N) by the
15 use of a viewing application 142. The viewing application 142 can be any suitable application such as a presentation program, a web browser, a media player, etc. The viewing application 142 may also be a web-based application. In one embodiment, the viewing application 142 may be a virtual assistant that uses voice instead of a visual representation to convey chunk contents, e.g., facilitating text-to-speech (TTS). In some embodiments, the
20 viewing application 142 may be an augmented reality, mixed reality, or virtual reality device and the chunk can be rendered within a virtual reality display. For example, a chunk may represent sensor data that can be updated in the background. A chunk may represent, for example, a virtual car configured in a shopping scenario that is updated on the fly and displayed in a virtual city.

[0039] It should be appreciated that the subject matter described herein may be implemented as a computer-controlled apparatus, a computer process, a computing system, or as an article of manufacture such as a computer-readable storage medium. Among many other benefits, the techniques shown herein improve efficiencies with respect to a wide range of computing resources. For instance, human interaction with a device may be
25 improved, as the use of the techniques disclosed herein enable a user to view and edit data from a wide range of file types while operating in one application to compose a chunk that includes data from the files. In addition, improved human interaction improves other computing resources such as processor and network resources, e.g., users can work from a reduced number of applications and reduce a user's computer interaction, reduce the
30

chances of an inadvertent input, reduce network traffic, and reduce computational cycles. The techniques disclosed herein reduce the need to download, start, maintain updates for, and toggle between, a number of applications, including a specialized presentation program. Also, instead of requiring users to copy and paste selected content into a specialized publication or presentation program, existing computer designs are improved in that all applications are synchronized and ultimately configured to improve the efficiency of a user's interactions. Other technical effects other than those mentioned herein can also be realized from implementations of the technologies disclosed herein.

[0040] In one embodiment, a storage service may be provided that is configured to manage and store chunks. In some embodiments, chunks may be implemented using tables of data. In one implementation, the data may have columns such as ID, Title, Author, and TimeDate. Each row of the table may be an instance of a chunk of information. For example, if Bob Smith is jotting ideas to obtain feedback on his sales and operations team, he may create three chunks called "Confirm receipt of goods," "Happy with sales process?" and "Likely to buy more?" The three chunks may be associated with the name Bob Smith, a timestamp of the time and date that the chunks were created, and each chunk may have different unique IDs. FIG. 2 illustrates one example table showing information pertaining to two chunks, each chunk having an ID which may conform to a globally unique identifier (GUID) scheme, a title for the chunk which may be a text string, an author of the chunk which may be a text string or some other identifier for a user, and a date/time that the chunk was modified which can be in a date and time format.

[0041] In some embodiments, additional or different columns may be included. Examples of columns may include location, type, or purpose. Additionally, columns may contain a file such as an image, video, document fragment, or an XML/HTML file. In some embodiments, a chunk of information may contain multi-dimensional data. For example, a rating column may be included which may in turn contain sub-columns for counts of 1 star, 2 stars, etc.

[0042] Multiple applications or services may read or write the data in a chunk, and allow for modifications to the schema of the data. The term "application" may refer to any application, website, add-in, SaaS service, etc. The storage service may receive information from one or more applications that are configured to receive user or system input pertaining to chunks. For example, inputs may be received via jotted notes, photos, automated insights from data analysis, insights generated from IoT sensors, etc.

[0043] The storage service may be remotely implemented such as on a server, or may be implemented on one or more devices. In some embodiments, the device providing the

storage service may be separate from the device where the application is executing. The application may read and/or write data to the storage service over a network. In some implementations, the schema of the data may include a column containing a list of who has access rights. The access rights may indicate, for example, that only an authenticated user such as BobSmith@Contoso.com and FredJacoby@Contoso.com may access the data. APIs may also be exposed to allow users to request or retrieve relevant chunks, such as those that the users have access to or are engaged with because of a shared task or project.

[0044] Chunks may be created and edited using any application that is capable of receiving user inputs and processing the inputs to generate or modify information pertaining to a unit of information such as a file. In one example, chunks may be edited using a text editor such as a Notes application on a smartphone. The chunks may also be edited using any number of more complex applications such as Microsoft Excel. In an embodiment, the chunks may be saved using a relatively simple form in the storage service, even when modified using a more complex editing application.

[0045] In an embodiment, an application may be configured to allow a user to insert one of the chunks into a document file. For example, if Bob wished to create a PowerPoint presentation entitled “Measuring customer satisfaction,” Bob may have an opening slide with that title, and a second slide called “Proposed Customer Survey.” Bob may then indicate to the application, for example by clicking on a button in the user interface, that he wishes to re-use a chunk of information that is accessible to Bob. In one embodiment, the application may provide a list of chunks that are available for insertion. In some embodiments, the list may be sorted by parameters such as time, most recent, or may provide an option to filter the list of chunks based on a search query term, date filtering, type, or other parameters.

[0046] Some embodiments may account for the scenario where not all chunk types may work in all document types. For example, a relatively simple text editor may support chunks along with basic text editing, and in this scenario, charts or pictures may not be provided as options for insertion. In some embodiments, the application may filter to those chunk types that the application supports. For example, if a chunk that supports multiple visualizations (e.g., text or graphs), then the chunk may be allowed to be inserted, but the available visualizations may be restricted to those that are supported. In some embodiments, chunks may be accessible to users who have a subscription to one or more chunks, or who have paid a fee to access the chunk. For example, an online store may provide chunks and/or associated visualization logic for a fee. In some embodiments, an internal online store or

repository may be provided for users within an organization or other entity. In scenarios that are internal to organizations where there can be a mix of internal and external chunks, the chunks may be subject to administrative controls provided by the organization. In some embodiments a chunk may have a licensing mechanism associated with the chunk. The application may check for licensing requirements, for example when attempting to refresh the document. The application may also check for rights-based access limitations as well as other restrictions.

[0047] In one embodiment, a column may be included for a chunk that specifies the types of content that are supported, such as text, images, vector drawings, etc. In some embodiments, alternate views of the same chunk data may be provided. For example, a chunk that is an image format may be described by words or sound in addition to being rendered as an image. In some embodiments, the type of a data chunk may be strongly defined. For example, 'PLAIN-TEXT' may be a simple type, and 'AVERAGE-RATING' could be a more complex type. 'PIE CHART' could be a subtype of 'CHART'. By providing for strongly defined types, visualizations can be specialized. For example, a complex type of a chart can be rendered such as an annotated 3-D chart with statistically significant deviations highlighted, but applications may be allowed to fall back to simpler visualizations (e.g., a plain bar chart), which can be enabled by defining '3D DEVIATION CHART' as a subtype of 'CHART'. Thus, when a user wants to change a visualization, the user can initially be shown relevant subtypes, parent types, or other related types.

[0048] Continuing with the example, Bob may choose to insert the two chunks "Happy with sales process?" and "Likely to buy more?" into a customer survey slide. In further embodiments, the application may allow a user to update a given chunk and the storage service may record a snapshot of the data as the data changes. In the example shown in FIG. 3, columns are included to indicate versioning and access rights. In some embodiments, versioning may be implemented in non-linear ways. For example, different entities such as users, bots, companies, and organizations may want different versions of a chunk. In some embodiments, chunks may be versioned using a tree-like or graph-like structure where chunks may spawn multiple versions in parallel. Different branches/edges may represent different edits in parallel, leading to disconnected versions or nodes of the same chunk. Every version or node could be considered the "latest" in its representative sub-tree/sub-graph. A management tool may be provided to update chunks on individual sub-trees/sub-graphs only, e.g., when determining licensing information. The management tool may also allow for merging information in two different sub-trees/sub-graphs and creating a new node

with the merged content. Furthermore, the management tool may allow for overriding a node with contents from a previous node. When a chunk is made up of other chunks, the management tool may be used to replace a chunk within the chunk with a different chunk or remove the chunk. Every edit of a chunk may be a new version. Each chunk version may be uniquely identified with an ID. In one embodiment, the differences between chunks may be stored instead of or in addition to the chunks themselves. This may facilitate tracking and checking the differences between versions, as well as save storage space. In other embodiments, multiple complete copies of various chunk versions may be stored. In one embodiment, a version may have a pointer to the previous version.

[0049] In some embodiments, a mechanism may be provided to address situations where conflicting or simultaneous edits by different authors are made to a chunk. In one embodiment, the prevailing edit may be determined as the last modified edit. In other embodiments, the conflicting or simultaneous edits may each be memorialized, tracked, and made available. For example, if a user in a marketing function and a user a research department each modified a list, then each edit may have potential usability by different end users. Thus, a mechanism may be provided where a user may specify the attributes that can determine which edits will be provided to that user. For example, the user may want to see the last edits by someone in marketing, or they want to see the last edits by a set of named individuals, or a central administrator may define rules for selecting edits which cannot be changed by end users. Other ways may be provided to determine how conflicting or simultaneous edits may be selected.

[0050] In some embodiments, a mechanism may be provided that allows end users to be associated with attributes that determine which updates and versions of updates are accessible to a given end user and/or for a given application.

[0051] Information may be saved in the document file to connect it to the chunks of information. For example, the document file may include the following information:

```
<powerpoint>
  <slide slide_id=1>
    ...
  </slide>
  <slide slide_id=2>
    <title>
      <format font=72>
        <p>Proposed Customer Survey</p>
```

```

                    </format>
                </title>
                <body>
                    <format bullets=true>
5                <p><information_chunk id=1273893 chunk_version=1.0>Happy with Sales
Process?</information_chunk></p>
                <p><information_chunk id=1273894 chunk_version=1.0>Likely to buy
more?</information_chunk></p>
                </format>
10            </body>
            </slide>
        </powerpoint>

```

[0052] In the example illustrated above, the second slide now contains a reference to the chunk storage service. For example, the tag “<information_chunk id=1273893
15 chunk_version=1.0>Happy with Sales Process?</information_chunk>” indicates the chunk identifier and the version. Additionally, the tags refer to two older chunks of data (version 1.0) which will then be retrieved from the chunk storage service and inserted into the document. In one embodiment, the location of the storage service may be assumed to be known to the applications that attempt to open the document. For example, Microsoft Office
20 applications may have built-in knowledge of the URL or other link to the storage service for chunks that are available to Microsoft Office applications. Additionally, Microsoft Office applications may have built-in knowledge of authentication methods for accessing the storage service. In other embodiments, the source may be explicitly stated as a property within the document file.

[0053] In some embodiments, the storage service may store groups of items which were originally authored to be part of a coherent list. Examples may include a sequential set of
25 photos, a set of tasks all of which are placed on the same board, and a set of bullets from a presentation slide. The ability to group items in sets may be beneficial when reusing information in chunks, when the information is to be used as a set. Additionally, users of
30 the sets of chunks may benefit by being notified when new items are added to a set, or when existing items in a set are deleted.

[0054] In one embodiment, a grouping mechanism may be implemented in the storage service by having each row of information include a shared attribute value. For example, the shared attribute may be the parent location URL as shown in the example table in FIG.

4. The shared attribute may also be indicative of an order, for example, of which item in the list is part of a given chunk.

[0055] In some embodiments, the level of granularity of the chunks may vary. For example, each row of a list of chunks in the storage service may be a full list or a full table, and methods may be provided for the author (or other user) to retrieve certain items from each chunk. While a chunk may be retrievable, the reference in the document where the chunk is inserted may refer to chunk items rather than the chunks themselves, as shown in the example below.

<body>

```

10      <format bullets=true>
        <information_chunk id=1273893 chunk_version=1.0 items=all>
        <p>Happy with Sales Process?</p>
        <p>Likely to buy more?</information_chunk></p>
        </information_chunk>
15      </format>

```

</body>

[0056] In an embodiment, applications which can process the document file may be able to perform operations when re-opening the document. For example, an application may be running on a computer device which is unable to connect over the network to the storage service (which may be referred to as “offline support”). In this case, the application may still be able to read the static copy of the information from the document file so that the application can display the version of the chunk as saved in the document. In one example, the reason for the storage service being unavailable is that a different user may have opened the application who has not been granted access to the storage service. In some

20

embodiments, the application may provide a notification to the user, for example by rendering a notification bar which displays “This document contains live information chunks. To view the latest version, click here to request access.” In response to clicking the request, the user may be taken to a different application or webpage where they can request access control to the storage service.

[0057] In one embodiment, if a user does not have access to the online data, then restrictions may be implemented to restrict allowing the unauthorized user from viewing the data offline as well. In some embodiments, access rights to the data may be included as one of the items optionally stored in the chunk definition. The access rights may be implemented to prevent the showing of the content if the access is insufficient. In some implementations,

IRM or DRM may be applied to the inline content of the chunk to ensure the user doesn't otherwise hack the document to access the content. In some embodiments, chunk contents may be encrypted. Additionally, IRM or DRM mechanisms may be inherited or transferred as chunks or portions of chunks are shared or forwarded. Other methods may be implemented for visual or audio files such as watermarking when chunks are printed or copied (e.g., screenshotted). In some embodiments, chunk content may not be included in the document so that without online access, the missing chunk content will result in having missing content in the document. For example, if there are concerns around access to the data, ensuring that the data is retrieved from the server by an authenticated user can ensure access controls are met. Additionally, a record of the access can be captured and stored.

[0058] When the application user does have access to the storage service or at least the desired chunk, the application may send a request to the storage service querying if a newer version of the chunk is available, or if items from the same location have been added or removed. In some embodiments, the application may provide the user with a prompt asking if they wish to update their document so that it contains all the latest versions of any inserted chunks. In other embodiments, the document may auto-refresh. In some embodiments, devices or applications may continuously or periodically poll the remote data store as to whether there are any updates. In one embodiment, the device or application may notify the storage service when it opens and closes the document and what chunks it contains so that the storage service can then send updates when they occur if the document is open. In still further embodiments, the user may explicitly set a property so that the document does not auto-refresh for other users or based on other conditions. For example, if a published report references demographic data for a given quarter, it may not be desirable for a user to open the document a year later when the demographic data has changed considerably. The findings in the report (which may be multiple pages of textual discourse) may wrongly reference observations about the demographic data that were valid a year ago but are no longer valid with the most up-to-date (latest version) demographic data. In still further embodiments, textual analysis may be performed against the document to identify where the document references older versions of the chunks so that those parts of the document may be highlighted for revision by the user.

[0059] In an embodiment, an application may be configured to allow the user to directly edit the chunks within the application. Chunks that are edited within one application may be used across multiple applications and in multiple documents, even where the context in which the information is used may be different. In some embodiments, edits that were made

to a chunk using one application may be applied to other applications and documents. Features may be implemented to limit or prevent propagation of changes in scenarios where the changes are inapplicable. For example, if the chunk is a presentation slide that summarizes sales for a quarter (and is used in a plurality of presentations across the organization), the chunk may make sense within this context and updates may be propagated. In another example, if the chunk is a table containing survey results such as the percentage of respondents who agree with a set of questions, then the table may also make sense to re-use. If the difference between versions of the chunks is minor rewording of sentences, it may make sense to apply edits across documents where the chunk is used.

[0060] When chunks are edited, additional properties may be added as columns that are associated with the chunk. For example, as shown in the example in FIG. 5, Fred may turn the customer service bulleted ideas into a customer survey (note that author, modified, access rights, parent location are removed for brevity only).

[0061] In some cases, older documents may have chunks that are locked to specific versions so that some data does not update in documents that have locked chunks. The example in FIG. 5 shows that “Sales process satisfaction” includes two versions 1.0 and 1.1. This may be useful when documents are historical and intended to include information from a point in time that is not to be updated. When a more recent change is made to the chunk, the changes may not be applied to the locked data in the older or historical documents. In one embodiment, a mechanism may be implemented that is configured to receive updates to the chunk and to bypass documents that have incorporated the chunks but are “frozen.” In one embodiment, only data from external sources may be frozen. In another embodiment, a mechanism may be implemented to access previous versions for updating. In further embodiments, a mechanism may be provided to save an update to all old copies.

[0062] In one embodiment, a mechanism may be implemented that determines whether an update is significant based on one or more factors that can be defined by users of a chunk. Insignificant updates or updates that are otherwise deemed unnecessary to a particular instance of a chunk can be bypassed. For example, the user may be prompted to select whether an update should be propagated. Additionally, a mechanism may be implemented to allow updates to be disabled, for example where changing historical legal filings may be prohibited. In one embodiment, an application may show all usages of a chunk in any of its versions and allow for approval for updating of each “version/document”-pair that are used in the document.

[0063] In the example above, additional columns may be added to capture how users

respond to the survey. The type of information may evolve from a more simple or linear representation to one that is more complex. In some implementations, sub-types may be generated and included in the chunk, for example “survey/customer satisfaction” (in contrast to “survey/lunch options” or “survey/idea voting”). Thus, semantic annotations may be used to provide visualizations.

[0064] In some of the examples detailed herein, the original content of the chunk is a ‘topic’ and new content may be directed to a question. In other examples, the original content may be context-neutral, but the new list content may be context-specific. As such, a variety of heuristics may be used by the application or survey to determine whether to request explicit input from the user before applying updates or whether updates should be applied automatically. For example, heuristics may include but are not limited to: the character edit distance before and after, whether the type of chunk has changed, the quantity of increased items, whether a chunk is purely numeric, whether the change is an updated photograph file, whether a chunk is a presentation slide, whether a chunk has introduced significant quantity of new columns, whether the semantic type of the chunk is deemed similar to what is already being used in the document, whether a convertor exists to modify the latest chunk type (e.g. survey) into the type which is present in the document (e.g. simplex), whether the chunk consists solely of data retrieved from a connection to an external source, e.g. SQL database or website, and the like. In an embodiment, a mechanism may be provided in which third parties can sign and lock chunks and thus prevent them from being changed by others, potentially including the owner of the chunk. Additionally, metadata may be modified indicating that a chunk or its specific contents were changed.

[0065] In one embodiment, the following sequence for updating chunks may be implemented: when the user makes an update to a chunk in the context of one document (or one application), the data may be always saved in the storage. When a second document referencing the chunk or a second application reading data from the chunk is launched, or in the case of co-authored documents, when an update is detected, then a determination may be made whether to (1) update the chunk in the document to use the latest version of the chunk automatically, or (2) notify the user that a change has been made that they may wish to accept, and whether to accept or ignore the change.

[0066] Not all applications may support providing notifications to the user of updates to chunks. For example, a simple phone application which includes a notetaking function may not be configured to provide a notification platform. Additionally, not all applications may contain logic capable of determining which version of the chunk to use. In such cases, the

computational logic to determine which version of the chunk to provide to a particular device or application may be performed on a server which may be network-connected to the storage or co-located with it. This may allow the phone application and its interfaces to the storage service to remain relatively simple. For example, the application may simply perform tasks for reading/writing to a list. The complex tasks of transforming and handling updates may be performed on the server. The server may be configured to perform update processing in a variety of scenarios that are not limited to phone applications.

[0067] With chunks saved in a storage service independent of individual user devices, the chunks may be exposed to a broader set of applications and users. In some embodiments, users and applications can query the storage service for relevant chunks. For example, chunks created in the same parent location or created by the same author in a given timeframe may be queried and retrieved. In such cases, the input parameters may be provided by the application or user and the results may be filtered to direct matches. As the use of chunks increase, chunk retrieval mechanisms may be implemented that may incorporate search engine techniques, where queries such as “show me all recently authored or updated chunks containing the keyword ‘customer’” may rely on standard information retrieval (IR) practices to find those chunks which contain larger numbers of mentions of the keyword “customer.” As such, the query may return sorted results by assigning a sub-score for how recently authored a chunk is, and another sub-score for how many times the chunk contains “customer.” The two sub-scores may be combined (e.g., summed) to compute how relevant a given chunk is. Other scenarios such as how often a given chunk is re-used, and whether the chunk is re-used by known contacts or within an organization or team, may be additional factors used to rank query results for chunks, for example. More generally, the chunk storage and management service may be integrated with query engine (e.g., a search crawler for enterprise search scenarios, or a database-like query for programmatic queries). In some embodiments, the visualization logic associated with a chunk, or other functionality, can be configured to translate any type of rich metadata into plain text to facilitate free text searching. Additionally, in conjunction with searches for chunks, a user may be provided the ability to select a chunk and be provided a list of similar chunks. In some embodiments, the list may be biased using a determined relevance based on social network data, previous searches, the current application, and other factors. Chunks may be indexed to facilitate being subjected to searches.

[0068] In various embodiments, a mechanism may be provided to render chunks of information. The chunks may be rendered within the context of an application, or

independently of an application. For example, chunks from the chunk storage service may be queried and rendered for review, or a chunk may be reviewed within an application.

[0069] In some embodiments, applications may be configured to render the chunk of information. The chunk may be rendered based on the type of information that it contains, as well as properties of the chunk as indicated by its metadata, which can be used to determine how the information is rendered. Some of the metadata associated with a chunk (e.g., author, time date, version, parent location, item, etc.) may not affect how the information is rendered in a document. Other properties such as the title or an associated image file may be used for rendering. For example, the application may render text as a plain text or may render an image within the document, wherever the user inserted the chunk in a document. If the original application (e.g., a phone note-taking application) is relaunched, then the original application may directly use the chunk(s) and not have any document content surrounding the chunk. The application may also render the chunk as text with metadata. In some embodiments, a mechanism or interface for changing between versions may be provided.

[0070] In some embodiments, a visualization mechanism for viewing, editing, and otherwise accessing chunks may be incorporated into an application. Alternatively or additionally, the visualization mechanism may be implemented as a standalone application or process. In some embodiments, the chunks may include or be linked to logic that enables an application to render the chunk information without the need for pre-existing functionality in the application or locally on the user device. In some embodiments, the logic may be configured to allow visualization of the chunk content. In one embodiment, the logic can provide a basic visualization for applications that are not configured to otherwise visualize the content of the chunk.

[0071] The associated chunk visualization logic may be implemented as code that can run on the user device or at the storage service when visualization is needed. For example, if the data in the chunk is a table, and the local device is attempting to visualize the table using Notepad, then the table may be rendered using logic provided for the chunk. The logic may be co-located with the chunk at the storage service, or available for download to the user device. Otherwise, the Notepad application may render the table using its limited rendering capabilities. For example, the contents of the table may be rendered as a list with delimiters.

[0072] In order to facilitate the implementation of chunks in various applications, the applications may conform to a standard interface for enabling the various techniques described herein. The standard interface may include a set of rules and interfaces that each

application must adhere to in order to be able to use a chunk regardless of its originating application and context. In some embodiments, add-ins that incorporate visualization and chunk management functionality may be implemented and made available to enable applications to use chunks within their documents.

5 **[0073]** In some embodiments, applications may override the base representation provided by the associated logic for a chunk. For example, the application may invoke another application to render the chunk data, such as a browser. The browser may then be used to render the chunk information according to the browser's capabilities. The application that is using the chunk may thus rely upon the browser or other application to render the chunk
10 for visualization by the user.

[0074] The logic may provide various capabilities, such as versioning and support for editing chunks. Furthermore, the logic may provide a user with the ability to visualize and provide input regarding the usage of the chunk. For example, when the originator of a chunk changes the data in the chunk, the current user may be provided a notification of the change,
15 and further may be provided a choice whether to incorporate the change. Additionally, the current user may be provided a notification of a change in visualization logic, and further may be provided a choice whether to update the visualization logic.

[0075] In some embodiments, the logic may be associated with the class that the chunk belongs to. For example, various classes of chunks may be defined, such as a "table_chunk"
20 class for chunks that have information in tabular form and is intended to be used in documents that include tables. The class and its associated logic may define a core representation for the class. In some embodiments, the logic may be stored with the chunks at the storage service. The logic may be sent along with the chunk when requested by an application. In other embodiments, the logic for various classes may be stored at another
25 location, for example at a master storage location for chunk visualization logic. In some embodiments, the logic for various chunk classes may be available via a URL. In one embodiment, users may define and submit their own chunk types through, for example, an online store or other centralized repository.

[0076] When the logic for various chunk classes are predefined and available for
30 download and use, when a chunk is generated, the logic does not need to be generated. The originating application only needs to generate the metadata for the instance, such as the title, the type of information in the chunk, who can access the chunk, the author, time/date stamp, class, and how to render the class (e.g., where to obtain the logic). In some embodiments, chunk-specific logic may be available in addition to class-specific logic. For example, the

class-specific logic for a table may be overridden for a particular table. In one embodiment, this can be achieved by applying the base display logic and layering on the custom logic, or by only applying the custom logic. There may be some applications that do not support the custom logic, in which case the logic can fall back to the default logic for that chunk type or class. The custom logic may also provide other visualization options such as those described herein. In some embodiments, multiple sets of logic may be associated with a chunk specifying different visualizations or different ways of editing.

[0077] In some embodiments, users may be presented with a number of available visualization options. Each visualization option may have their own logic and can process the data for the chunk and generate a rendering. The types of renderings can range from pure text, pure graphics, and any other type or combination of renderings. In some embodiments, purely graphic oriented applications may be able to render a variety of chunk types. For example, a Paint application may be able to render a table graphically, even if editing capabilities may not be provided. In some embodiments, the visualization logic may define how to render chunks if other chunks are also present in the document.

[0078] In some embodiments, a ranked list of visualizations may be provided, for example a list of the most popular visualization options. The list may be provided as a text rendering. In another example, a graphical output may be generated that shows an animation depicting a rotating carousel with the most popular visualizations. In one embodiment, each visualization in the list may have a preview of the type of rendering that will be provided with that visualization, which the user may select.

[0079] In an embodiment, the list may be selected or prioritized based on one or more factors including:

[0080] Which visualization was last used to render the chunk

[0081] What is the strong type of the chunk

[0082] Which visualizations are related to that type (fully or partially)

[0083] Which visualizations is the user licensed for.

[0084] In some embodiments, the user may be provided the capability to filter the visualizations based on capabilities (e.g., edit vs. read only, black & white vs. color theming).

[0085] In some scenarios, the only available visualizations in an application (e.g., Notepad) may be able to facilitate the rendering of information, but not its editing. For example, a star rating can be shown, but since its data comprises a list of all previous submitted ratings with the total being averaged, there may be no way for a user to edit that

data in Notepad without losing the underlying ratings data. In such scenarios, an application may notify the user that a given chunk of data is not editable by the user or that editing will result in data loss. Alternatively, the application may not allow the chunk to be edited without an explicit notification (e.g., editing functions may be disabled). FIGURE 6 illustrates a block diagram 600 for identifying chunks, storing chunks, and providing the chunks to an application. In some embodiments, chunk sources 601 may include documents, emails, spreadsheets, presentations, graphics files, and the like. These documents may be accessed via a network by chunk identification engine 602. Chunk identification engine 602 may receive chunks, including metadata associated with the chunks, to associate chunks with metadata and related rendering logic. In some embodiments, chunks may be stored in chunk store 604, which may be implemented by a storage service as discussed herein.

[0086] In some embodiments, chunks stored in chunk store 604 may be retrieved by software applications, such as presentation software 606 or word processing software 608. These applications may be configured to include authoring panes 610 and 614 in which main document content is displayed. Applications 606 and 608 may also include chunk add-ins 612 and 616. In some embodiments, the chunk add-ins may receive chunks and insert portions of the chunk information in documents being edited by applications 606 and 608 and displayed in authoring panes 610 and 614. In other embodiments, the functionality provided by add-ins may be included in the functionality of the applications.

[0087] FIGURE 7 illustrates a block diagram 700 for generating, storing, and re-using chunks in conjunction with a team or organization 706. In some embodiments, chunk identification engine 602 receives chunks to be processed from collaboration server 716, email server 714, cloud document server 716, or the like. In some embodiments, the chunks stored on these servers are associated with organization 706. This enables chunks to be identified that are particular to organization 706.

[0088] For example, collaboration server 716 may store slide presentations, spreadsheets, word processing documents, emails, calendar appointments, or the like. In some embodiments, chunk identification engine 602 may have access to retrieve the documents stored in collaboration server 716. Additionally or alternatively, collaboration server 716 may provide stored documents to chunk identification engine 602. Chunk identification engine 602 may also be implemented as an add-in executing within collaboration server 716.

[0089] Chunk identification engine 602 may similarly receive artifacts from email server 714. These artifacts may include emails, attachments, or any other type of document.

[0090] In some embodiments, while cloud document server 716 exists physically outside of organization 706, the documents contained thereon may still be considered part of the organization.

5 [0091] FIGURE 8 illustrates a block diagram 800 showing how a chunk is updated by various users and updates provided to users. In some embodiments, chunk identification engine 602 is configured to track chunk 850 that is inserted in document 801, document 802, document 804 and document 807. These documents have associated users 808, 810, 812, and 816.

10 [0092] In some embodiments, as document 802, document 804, and document 806 are being updated, when their respective inserted chunks 850 are updated, the updates may be received by chunk identification engine 602 and document 802, document 804, and document 806 may each be notified of any updates to the chunk and/or its associated visualization logic. In this example, document 804 associated with user 810 may not currently be in use. However, chunk identification engine 602 may track the updates to
15 chunk 850 and when document 804 is later opened, chunk identification engine 602 may send the document 804 the accumulated updates.

[0093] FIGURE 9 shows a document-based application 900 including a chunk visualization pane 940 for chunks that are inserted in the document. In one embodiment, application 900 executes on computing device 902, and includes a document-based
20 application including a menu 910, a document authoring pane 990, and a chunk visualization pane 940.

[0094] Chunk visualization pane 940 lists three chunks found in the document: chunk 1 942, chunk 2 949, and chunk 3 952, with arrows associating the chunks with one or more instances of the chunks being used in the document.

25 [0095] FIGURE 10 illustrates aspects of a routine 1000 for enabling aspects of the techniques disclosed herein as shown and described below. It should be understood that the operations of the methods disclosed herein are not presented in any particular order and that performance of some or all of the operations in an alternative order(s) is possible and is contemplated. The operations have been presented in the demonstrated order for ease of
30 description and illustration. Operations may be added, omitted, and/or performed simultaneously, without departing from the scope of the appended claims.

[0096] It also should be understood that the illustrated methods can end at any time and need not be performed in their entirety. Some or all operations of the methods, and/or substantially equivalent operations, can be performed by execution of computer-readable

instructions included on a computer-storage media, as defined below. The term “computer-readable instructions,” and variants thereof, as used in the description and claims, is used expansively herein to include routines, applications, application modules, program modules, programs, components, data structures, algorithms, and the like. Computer-readable instructions can be implemented on various system configurations, including single-processor or multiprocessor systems, minicomputers, mainframe computers, personal computers, hand-held computing devices, microprocessor-based, programmable consumer electronics, combinations thereof, and the like.

[0097] Thus, it should be appreciated that the logical operations described herein are implemented (1) as a sequence of computer implemented acts or program modules running on a computing system and/or (2) as interconnected machine logic circuits or circuit modules within the computing system. The implementation is a matter of choice dependent on the performance and other requirements of the computing system. Accordingly, the logical operations described herein are referred to variously as states, operations, structural devices, acts, or modules. These operations, structural devices, acts, and modules may be implemented in software, in firmware, in special purpose digital logic, and any combination thereof.

[0098] For example, the operations of the routine 900 are described herein as being implemented, at least in part, by modules running the features disclosed herein and can be a dynamically linked library (DLL), a statically linked library, functionality produced by an application programming interface (API), a compiled program, an interpreted program, a script or any other executable set of instructions. Data can be stored in a data structure in one or more memory components. Data can be retrieved from the data structure by addressing links or references to the data structure.

[0099] Although the following illustration refers to the components of the figures, it can be appreciated that the operations of the routine 1000 may be also implemented in many other ways. For example, the routine 1000 may be implemented, at least in part, by a processor of another remote computer or a local circuit. In addition, one or more of the operations of the routine 1000 may alternatively or additionally be implemented, at least in part, by a chipset working alone or in conjunction with other software modules. In the example described below, one or more modules of a computing system can receive and/or process the data disclosed herein. Any service, circuit or application suitable for providing the techniques disclosed herein can be used in operations described herein.

[00100] FIGURE 10 illustrates aspects of a routine 1000 for enabling aspects of the

techniques disclosed herein as shown and described below. The operations in FIG. 10 can be performed, for example, by the computing device 1300 of FIG. 13, as described above with respect to any one of FIGS. 1-9.

5 [00101] At operation 1001, data indicative of a persistent data object may be received from a data store by a computing device. In an embodiment, the persistent data object may be associated with a class indicative of a data type for information contained in the persistent data object. Additionally, the persistent data object may be operable to be inserted in a file generated by an application executing on the computing device.

10 [00102] At operation 1003, visualization logic associated with the class may be received by the computing device.

[00103] At operation 1004, based at least in part on the visualization logic, a method of visualization of the persistent data object is selected.

15 [00104] At operation 1005, the information in the persistent data object may be rendered using the selected method of visualization. The persistent data object may be rendered using the visualization logic based on other factors, or based on a user selection. For example, the visualization logic might provide a more modern or more appropriate visualization.

[00105] At operation 1007, other information in the document may be rendered using native rendering capabilities of the application.

20 [00106] In some embodiments, the visualization logic may be configured to render, on a display device, one or more visualization options for rendering the information contained in the persistent data object.

[00107] In some embodiments, the visualization logic may be configured to render an interface for editing the information contained in the persistent data object.

[00108] In some embodiments, the persistent data object comprises a strongly-defined type.

25 [00109] In some embodiments, the interface displays different revisions of the persistent data object and differences between the revisions.

[00110] In some embodiments, the visualization logic may be configured to render an interface for selecting an option to use the native rendering capabilities of the application to render the information. The interface may further display different revisions of the persistent data object

30 [00111] In some embodiments, the persistent data object may be further associated with one or more of an author, date/time stamp, version, and title. Additionally, the visualization logic may be operable to render the author, date/time stamp, version, and title.

[00112] In some embodiments, the visualization logic may be operable to render an

interface for authorizing users who are allowed to view or edit the persistent data object.

[00113] In some embodiments, the persistent data object and its associated metadata may be stored as a table. The visualization logic may be operable to render information contained in the persistent data object and its associated metadata in a tabular format.

5 [00114] In some embodiments, the visualization logic is operable to render one or more of:

[00115] sources of chunk information,

[00116] authors of parts of the persistent data object;

[00117] an option to merge two or more chunks;

[00118] visualizations based on themes;

10 [00119] visualizations based on text- to- speech;

[00120] time of update; and

[00121] author information (e.g., job title, division, profile photo, etc.).

[00122] Other information may also be rendered if available.

15 [00123] In some embodiments, the visualization options may comprise a list of rendering options. Additionally, the list may be prioritized or ranked. The ranking may be based on visualizations that are semantically or functionally relevant visualizations given a type of the persistent data object.

[00124] In some embodiments, the visualization logic may be co-located in the data store.

20 In some embodiments, the visualization logic may be stored remotely from the data store and is accessed via a link to a repository for visualization logic, the repository storing a plurality of types of persistent data objects.

[00125] In some embodiments, the persistent data object and an associated visualization is associated with a hierarchy of strong types.

25 [00126] In some embodiments, a given persistent data object may be locked from editing, and an associated visualization logic may not be allowed to edit the persistent data object. Additionally, the application may not natively permit editing of the persistent data object in this case. Furthermore, a notification may be rendered that the information cannot be edited.

[00127] In some embodiments, the visualization logic is filtered based on whether a user is licensed to access the visualization logic.

30 [00128] FIGURE 11 illustrates aspects of a routine 1100 for enabling aspects of the techniques disclosed herein as shown and described below. The operations in FIG. 11 can be performed, for example, by the computing device 1300 of FIG. 13, as described above with respect to any one of FIGS. 1-9.

[00129] At operation 1101, a persistent data object may be received. In an embodiment, the

persistent data object may be usable by an application executing on the computing device to insert data in a document being edited by the application based on information contained in the persistent data object.

5 [00130] At operation 1103, visualization logic associated the persistent data object is accessed.

[00131] At operation 1104, a method of visualization of the persistent data object is selected.

[00132] At operation 1105, the selected method of visualization is used to render the information in the persistent data object.

10 [00133] At operation 1107, the native capabilities of the application are used to render other information in the document.

[00134] In some embodiments, the persistent data object may be originally generated by another application. The persistent data object may be received from a remote data store configured to store a plurality of persistent data objects.

15 [00135] In some embodiments, the application may be configured to optionally render the information in the persistent data object by overriding the rendering provided by the visualization logic.

[00136] In some embodiments, the visualization logic may be predetermined for a class associated with the persistent data object.

20 [00137] In some embodiments, the visualization logic may be configured to render an interface operable to provide rendering options for the information contained in the persistent data object.

[00138] In some embodiments, an indication is rendered that at least some properties of the persistent data object are incompatible with available visualizations or with the selected visualization.

25 [00139] In some embodiments, the persistent data object comprises a strongly-defined type

[00140] FIGURE 12 illustrates aspects of a routine 1200 for enabling aspects of the techniques disclosed herein as shown and described below. The operations in FIG. 12 can be performed, for example, by the computing device 1300 of FIG. 13, as described above with respect to any one of FIGS. 1-9.

30 [00141] At operation 1201, data indicative of a persistent data object may be stored in a data store. In some embodiments, the persistent data object may be associated with a class indicative of a data type for information contained in the persistent data object. Additionally, the persistent data object may be operable to be inserted in a file generated by an application.

[00142] At operation 1203, in response to a request from an application for the persistent

data object, visualization logic associated with the class is accessed.

[00143] At operation 1205, the persistent data object and the visualization logic may be sent to a computing device executing the application.

5 [00144] In an embodiment, the visualization logic is operable to enable the computing device to render the information in the persistent data object using the application. In some embodiments, the visualization logic may be used to render the information even if the application is capable of rendering the information.

10 [00145] In an embodiment, the visualization logic may be operable to enable the computing device to filter the information in the persistent data object based on one or more metadata associated with the persistent data object. The visualization logic may further be operable to replace data in the persistent data object with data of the persistent data object's metadata. For example, translations may be stored in the metadata.

15 [00146] In an embodiment, the visualization logic may be operable to render an interface configured to allow a user to add and modify metadata associated with the persistent data object.

[00147] In an embodiment, the visualization logic may be operable to receive and render a notification that the persistent data object has an available update.

20 [00148] In an embodiment, the visualization logic may be operable to present a listing of available visualizations for the persistent data object. The available visualizations may have associated logic for processing some or all of the information in the persistent data object and generating a rendering.

25 [00149] FIGURE 13 shows additional details of an example computer architecture 1300 for a computer, such as a computing device executing chunk management function 1380 and chunk visualization logic 1390 (FIGURE 13), capable of executing the program components described herein. Thus, the computer architecture 1300 illustrated in FIGURE 13 illustrates an architecture for a server computer, a mobile phone, a PDA, a smart phone, a desktop computer, a netbook computer, a tablet computer, and/or a laptop computer. The computer architecture 1300 may be utilized to execute any aspects of the software components presented herein.

30 [00150] The computer architecture 1300 illustrated in FIGURE 13 includes a central processing unit 1302 ("CPU"), a system memory 1304, including a random access memory 13013 ("RAM") and a read-only memory ("ROM") 13013, and a system bus 1310 that couples the memory 1304 to the CPU 1302. A basic input/output system containing the basic routines that help to transfer information between elements within the computer

architecture 1300, such as during startup, is stored in the ROM 1306. The computer architecture 1300 further includes a mass storage device 1312 for storing an operating system 1307. Mass storage device 1312 may further include chunk visualization logic 1390 and chunk management function 1380, which include some or all of the aspects of chunk visualization functionality as disclosed herein and chunk management functionality as disclosed herein.

[00151] The mass storage device 1312 is connected to the CPU 1302 through a mass storage controller (not shown) connected to the bus 1310. The mass storage device 1312 and its associated computer-readable media provide non-volatile storage for the computer architecture 1300. Although the description of computer-readable media contained herein refers to a mass storage device, such as a solid state drive, a hard disk or CD-ROM drive, it should be appreciated by those skilled in the art that computer-readable media can be any available computer storage media or communication media that can be accessed by the computer architecture 1300.

[00152] Communication media includes computer readable instructions, data structures, program modules, or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any delivery media. The term “modulated data signal” means a signal that has one or more of its characteristics changed or set in a manner so as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer-readable media.

[00153] By way of example, and not limitation, computer storage media may include volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules or other data. For example, computer media includes, but is not limited to, RAM, ROM, EPROM, EEPROM, flash memory or other solid-state memory technology, CD-ROM, digital versatile disks (“DVD”), HD-DVD, BLU-RAY, or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computer architecture 1300. For purposes of the claims, the phrase “computer storage medium,” “computer-readable storage medium” and variations thereof, does not include waves, signals, and/or other transitory and/or intangible communication media, per se.

[00154] According to various configurations, the computer architecture 1300 may operate in a networked environment using logical connections to remote computers through the network 13513 and/or another network (not shown). The computer architecture 1300 may connect to the network 13513 through a network interface unit 1314 connected to the bus 1310. It should be appreciated that the network interface unit 1314 also may be utilized to connect to other types of networks and remote computer systems. The computer architecture 1300 also may include an input/output controller 13113 for receiving and processing input from a number of other devices, including a keyboard, mouse, or electronic stylus (not shown in FIGURE 13). Similarly, the input/output controller 13113 may provide output to a display screen, a printer, or other type of output device (also not shown in FIGURE 13).

[00155] It should be appreciated that the software components described herein may, when loaded into the CPU 1302 and executed, transform the CPU 1302 and the overall computer architecture 1300 from a general-purpose computing system into a special-purpose computing system customized to facilitate the functionality presented herein. The CPU 1302 may be constructed from any number of transistors or other discrete circuit elements, which may individually or collectively assume any number of states. More specifically, the CPU 1302 may operate as a finite-state machine, in response to executable instructions contained within the software modules disclosed herein. These computer-executable instructions may transform the CPU 1302 by specifying how the CPU 1302 transitions between states, thereby transforming the transistors or other discrete hardware elements constituting the CPU 1302.

[00156] Encoding the software modules presented herein also may transform the physical structure of the computer-readable media presented herein. The specific transformation of physical structure may depend on various factors, in different implementations of this description. Examples of such factors may include, but are not limited to, the technology used to implement the computer-readable media, whether the computer-readable media is characterized as primary or secondary storage, and the like. For example, if the computer-readable media is implemented as semiconductor-based memory, the software disclosed herein may be encoded on the computer-readable media by transforming the physical state of the semiconductor memory. For example, the software may transform the state of transistors, capacitors, and/or other discrete circuit elements constituting the semiconductor memory. The software also may transform the physical state of such components in order to store data thereupon.

[00157] As another example, the computer-readable media disclosed herein may be

implemented using magnetic or optical technology. In such implementations, the software presented herein may transform the physical state of magnetic or optical media, when the software is encoded therein. These transformations may include altering the magnetic characteristics of particular locations within given magnetic media. These transformations
5 also may include altering the physical features or characteristics of particular locations within given optical media, to change the optical characteristics of those locations. Other transformations of physical media are possible without departing from the scope and spirit of the present description, with the foregoing examples provided only to facilitate this discussion.

10 **[00158]** In light of the above, it should be appreciated that many types of physical transformations take place in the computer architecture 1300 in order to store and execute the software components presented herein. It also should be appreciated that the computer architecture 1300 may include other types of computing devices, including hand-held computers, embedded computer systems, personal digital assistants, and other types of
15 computing devices known to those skilled in the art. It is also contemplated that the computer architecture 1300 may not include all of the components shown in FIGURE 13, may include other components that are not explicitly shown in FIGURE 13, or may utilize an architecture completely different than that shown in FIGURE 13.

[00159] The disclosure presented herein may be considered in view of the following
20 example clauses.

[00160] Example Clause A, a method for rendering shared data, the method comprising:

receiving, from a data store by a computing device, data indicative of a
persistent data object, wherein:

the persistent data object is associated with a class indicative of a data type
25 for information contained in the persistent data object; and

the persistent data object is operable to be inserted in a file generated by an
application executing on the computing device;

receiving, by the computing device, visualization logic associated with the
class;

30 based at least in part on the visualization logic, selecting a method of
visualization of the persistent data object;

rendering the information in the persistent data object using the selected
method of visualization; and

rendering other information in the document using native rendering

capabilities of the application.

[00161] Example Clause B, the method of Example Clause A, wherein the visualization logic is configured to render, on a display device, one or more visualization options for rendering the information contained in the persistent data object.

5 [00162] Example Clause C, the method of Example Clause B, wherein the persistent data object comprises a strongly-defined type.

[00163] Example Clause D, the method of any one of Example Clauses A through C, wherein the visualization logic is configured to render an interface for selecting an option to use the native rendering capabilities of the application to render the information.

10 [00164] Example Clause E, the method of any one of Example Clauses A through D, wherein the interface displays different revisions of the persistent data object.

[00165] Example Clause F, the method of any one of Example Clauses A through D, wherein the visualization logic is operable to render an interface for authorizing users who are allowed to view or edit the persistent data object.

15 [00166] Example Clause G, the method of any one of Example Clauses A through F, wherein the visualization logic is operable to render one or more of:

sources of chunk information;

authors of parts of the persistent data object, an option to merge two or more chunks; visualizations based on themes;

20 visualizations based on text-to-speech;

time of update; or

author information.

[00167] Example Clause H, the method of any one of Example Clauses A through G, wherein the visualization options comprise a ranked list of rendering options, the ranking
25 based on visualizations that are semantically or functionally relevant visualizations given a type of the persistent data object.

[00168] Example Clause I, the method of any one of Example Clauses A through G, further comprising associating the persistent data object and an associated visualization with a hierarchy of strong types.

30 [00169] Example Clause J, the method of any one of Example Clauses A through I, wherein the visualization logic is filtered based on whether a user is licensed to access the visualization logic.

[00170] While Example Clauses A through J are described above with respect to a method, it is understood in the context of this disclosure that the subject matter of Example Clauses

A through J can additionally or alternatively be implemented by a device or via computer readable storage media.

[00171] Example Clause K, a computing device comprising:

one or more processors;

5 a memory in communication with the one or more processors, the memory having computer-readable instructions stored thereupon which, when executed by the one or more processors, cause the computing device to:

10 receive a persistent data object, wherein the persistent data object is usable by an application executing on the computing device to insert data in a document being edited by the application based on information contained in the persistent data object;

access visualization logic associated with the persistent data object;

based at least in part on the visualization logic, selecting a method of visualization of the persistent data object;

15 use the selected method of visualization to render the information in the persistent data object; and

use native rendering capabilities of the application to render other information in the document.

20 **[00172]** Example Clause L, the computing device of Example Clause K, wherein the persistent data object comprises a strongly-defined type.

[00173] Example Clause M, the computing device of Example Clause K or Example Clause L, wherein the visualization logic is filtered based on whether a user is licensed to access the visualization logic.

25 **[00174]** Example Clause N the computing device of Example Clause K through Example Clause M, wherein the visualization logic is predetermined for a class associated with the persistent data object.

30 **[00175]** Example Clause O, the computing device of Example Clause K through Example Clause N, further comprising rendering an indication that at least some properties of the persistent data object are incompatible with available visualizations or with the selected visualization.

[00176] While Example Clauses K through O are described above with respect to a computing device, it is understood in the context of this disclosure that the subject matter of Example Clauses K through O can additionally or alternatively be implemented by a method or via computer readable storage media.

[00177] Example Clause P, a method for managing shared data, the method comprising:

storing, in a data store, data indicative of a persistent data object, wherein:

the persistent data object is associated with a class indicative of a data type for information contained in the persistent data object; and

5 the persistent data object is operable to be inserted in a file generated by an application;

in response to a request from an application for the persistent data object, accessing visualization logic associated with the class; and

10 sending the persistent data object and the visualization logic to a computing device executing the application, wherein:

the visualization logic is operable to enable the computing device to render the information in the persistent data object and render other information using the application.

[00178] Example Clause Q, the method of Example Clause P, wherein the visualization logic is operable to enable the computing device to filter the information in the persistent data object based on one or more metadata associated with the persistent data object.

[00179] Example Clause R, the method of Example Clause P and Example Clause Q, wherein the visualization logic is operable to render an interface configured to allow a user to add and modify metadata associated with the persistent data object.

20 [00180] Example Clause S, the method of Example Clauses P through R, wherein the visualization logic is operable to receive and render a notification that the persistent data object has an available update.

[00181] Example Clause T, the method of Example Clauses P through S, wherein the visualization logic is operable to present a listing of available visualizations for the persistent data object, the available visualizations having associated logic for processing some or all of the information in the persistent data object and generating a rendering.

25 [00182] While Example Clauses P through T are described above with respect to a method, it is understood in the context of this disclosure that the subject matter of Example Clauses P through T can additionally or alternatively be implemented as a device or via computer readable storage media.

30 [00183] Although the techniques have been described in language specific to structural features and/or methodological acts, it is to be understood that the appended claims are not necessarily limited to the features or acts described. Rather, the features and acts are described as example implementations of such techniques.

[00184] The operations of the example processes are illustrated in individual blocks and summarized with reference to those blocks. The processes are illustrated as logical flows of blocks, each block of which can represent one or more operations that can be implemented in hardware, software, or a combination thereof. In the context of software, the operations represent computer-executable instructions stored on one or more computer-readable media that, when executed by one or more processors, enable the one or more processors to perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, modules, components, data structures, and the like that perform particular functions or implement particular abstract data types. The order in which the operations are described is not intended to be construed as a limitation, and any number of the described operations can be executed in any order, combined in any order, subdivided into multiple sub-operations, and/or executed in parallel to implement the described processes. The described processes can be performed by resources associated with one or more device(s) such as one or more internal or external CPUs or GPUs, and/or one or more pieces of hardware logic such as FPGAs, DSPs, or other types of accelerators.

[00185] All of the methods and processes described above may be embodied in, and fully automated via, software code modules executed by one or more general purpose computers or processors. The code modules may be stored in any type of computer-readable storage medium or other computer storage device. Some or all of the methods may alternatively be embodied in specialized computer hardware.

[00186] Conditional language such as, among others, "can," "could," "might" or "may," unless specifically stated otherwise, are understood within the context to present that certain examples include, while other examples do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that certain features, elements and/or steps are in any way required for one or more examples or that one or more examples necessarily include logic for deciding, with or without user input or prompting, whether certain features, elements and/or steps are included or are to be performed in any particular example. Conjunctive language such as the phrase "at least one of X, Y or Z," unless specifically stated otherwise, is to be understood to present that an item, term, etc. may be either X, Y, or Z, or a combination thereof.

[00187] Any routine descriptions, elements or blocks in the flow diagrams described herein and/or depicted in the attached figures should be understood as potentially representing modules, segments, or portions of code that include one or more executable instructions for implementing specific logical functions or elements in the routine. Alternate

implementations are included within the scope of the examples described herein in which elements or functions may be deleted, or executed out of order from that shown or discussed, including substantially synchronously or in reverse order, depending on the functionality involved as would be understood by those skilled in the art. It should be emphasized that

5 many variations and modifications may be made to the above-described examples, the elements of which are to be understood as being among other acceptable examples. All such modifications and variations are intended to be included herein within the scope of this disclosure and protected by the following claims.

CLAIMS

1. A method for rendering shared data, the method comprising:
 - receiving, from a data store by a computing device, data indicative of a persistent data object, wherein:
 - the persistent data object is associated with a class indicative of a data type for information contained in the persistent data object;
 - and
 - the persistent data object is operable to be inserted in a file generated by an application executing on the computing device;
 - receiving, by the computing device, visualization logic associated with the class;
 - based at least in part on the visualization logic, selecting a method of visualization of the persistent data object;
 - rendering the information in the persistent data object using the selected method of visualization; and
 - rendering other information in the document using native rendering capabilities of the application.
2. The method of claim 1, wherein the visualization logic is configured to render, on a display device, one or more visualization options for rendering the information contained in the persistent data object.
3. The method of claim 1, wherein the persistent data object comprises a strongly-defined type.
4. The method of claim 1, wherein the visualization logic is configured to render an interface for selecting an option to use the native rendering capabilities of the application to render the information.
5. The method of claims 4, wherein the interface displays different revisions of the persistent data object.
6. The method of claim 1, wherein the visualization logic is operable to render an interface for authorizing users who are allowed to view or edit the persistent data object.
7. The method of claim 1, wherein the visualization logic is operable to render one or more of:
 - sources of chunk information;
 - authors of parts of the persistent data object, an option to merge two or more chunks;
 - visualizations based on themes;

visualizations based on text-to-speech;
time of update; or
author information.

8. The method of claim 2, wherein the visualization options comprise a ranked list of rendering options, the ranking based on visualizations that are semantically or functionally relevant visualizations given a type of the persistent data object.

9. The method of claim 1, further comprising associating the persistent data object and an associated visualization with a hierarchy of strong types.

10. The method of claim 1, wherein the visualization logic is filtered based on whether a user is licensed to access the visualization logic.

11. A computing device comprising:
one or more processors;
a memory in communication with the one or more processors, the memory having computer-readable instructions stored thereupon which, when executed by the one or more processors, cause the computing device to:

receive a persistent data object, wherein the persistent data object is usable by an application executing on the computing device to insert data in a document being edited by the application based on information contained in the persistent data object;

access visualization logic associated with the persistent data object;
based at least in part on the visualization logic, selecting a method of visualization of the persistent data object;

use the selected method of visualization to render the information in the persistent data object; and

use native rendering capabilities of the application to render other information in the document.

12. The computing device of claim 11, wherein the visualization logic is predetermined for a class associated with the persistent data object.

13. The computing device of claim 11, further comprising rendering an indication that at least some properties of the persistent data object are incompatible with available visualizations or with the selected visualization.

14. A method for managing shared data, the method comprising:

storing, in a data store, data indicative of a persistent data object, wherein:
the persistent data object is associated with a class indicative

of a data type for information contained in the persistent data object;
and

the persistent data object is operable to be inserted in a file
generated by an application;

in response to a request from an application for the persistent data object,
accessing visualization logic associated with the class; and

sending the persistent data object and the visualization logic to a computing
device executing the application, wherein:

the visualization logic is operable to enable the computing device to render
the information in the persistent data object and render other information using the
application.

15. The method of claim 14, wherein the visualization logic is operable to
receive and render a notification that the persistent data object has an available update.




FIGURE 1

ID	Title (text)	Author (person)	Modified (timedate)
1273893	Happy with sales process?	Bob Smith	15/9/2018 18:31
1273894	Likely to buy more?	Bob Smith	15/9/2018 18:32

FIGURE 2

ID	Ver	Title (text)	Author (person)	Modified (timestamp)	Access Rights	Type
1273893	1.0	Happy with sales process?	Bob Smith	15/9/2018 18:31	Bob, Fred	Simplex
	1.1	Sales process satisfaction	Fred Jacoby	18/9/2018 10:12	Bob, Fred	Simplex
1273894	1.0	Likely to buy more?	Bob Smith	15/9/2018 18:32	Bob, Fred	Simplex
	1.1	Likelihood of additional purchase	Fred Jacoby	18/9/2018 10:12	Bob, Fred	Simplex

FIGURE 3

ID	Ver	Title (text)	Author (person)	Modified (timestamp)	Access Rights	Type	Parent Location (URL)	Item
1273893	1.0	Happy with sales process?	Bob Smith	15/9/2018 18:31	Bob, Fred	Simplex	Notes.contoso365.com/customerfeedback	1
	1.1	Sales process satisfaction	Fred Jacoby	18/9/2018 10:12	Bob, Fred	Simplex	Notes.contoso365.com/customerfeedback	1
1273894	1.0	Likely to buy more?	Bob Smith	15/9/2018 18:32	Bob, Fred	Simplex	Notes.contoso365.com/customerfeedback	2
	1.1	Additional Purchase Likelihood	Fred Jacoby	18/9/2018 10:12	Bob, Fred	Simplex	Notes.contoso365.com/customerfeedback	2
	1.2	Additional Purchase Likelihood	Fred Jacoby	18/9/2018 10:12	Bob, Fred	Simplex	Notes.contoso365.com/customerfeedback	3
1290761	1.0	Review budget	Bob Smith	22/10/2018 11:17	Bob	Simplex	Notes.contoso365.com/Octoberagenda	1
1291281	1.0	Frequency of historic purchase	Bob Smith	22/10/2018 15:13	Bob, Fred	Simplex	Notes.contoso365.com/customerfeedback/2	2

FIGURE 4


5/13

ID	Ver	Title (text)	Type	Item	Happy Count	Neutral Count	Unhappy Count	Average
1273893	1.0	Happy with sales process?	Simplex	1				
	1.0	Sales process satisfaction	Simplex	1				
	1.1	Sales process satisfaction	Simplex	1				
	1.2	How satisfied are you with our sales process?	Survey	1	0	0	0	null
1273894	1.0	Likely to buy more?	Simplex	2				
	1.1	Additional Purchase Likelihood	Simplex	2				
	1.2	Additional Purchase Likelihood	Simplex	3				
	1.2	How likely are you to purchase from us again?	Survey	3	0	0	0	Null
1290761	1.0	Review budget	Simplex	1				
1291281	1.0	Frequency of historic purchase	Simplex	2				
	1.1	How often have you purchased from us previously?	Survey	2	0	0	0	Null


FIGURE 5

6/13

600

**FIGURE 6**

600



ORGANIZATION 706

FIGURE 7

800

8/13




FIGURE 8




FIGURE 9

10/13

1000




FIGURE 10

11/13

1100




FIGURE 11

12/13

1200




FIGURE 12




FIGURE 13

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/038823

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/54 G06F8/33
ADD. G06Q10/10

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F G06Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	HECTOR GONZALEZ ET AL: "Google fusion tables", PROCEEDINGS OF THE 1ST ACM SYMPOSIUM ON CLOUD COMPUTING, SOCC '10, ACM PRESS, NEW YORK, NEW YORK, USA, 10 June 2010 (2010-06-10), pages 175-180, XP058190179, DOI: 10.1145/1807128.1807158 ISBN: 978-1-4503-0036-0 abstract figure 1 page 175, left-hand column, line 1 - page 176, left-hand column, line 35 page 178, left-hand column, line 28 - page 180, left-hand column, line 23 ----- -/-	1-15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

3 September 2019

Date of mailing of the international search report

11/09/2019

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Wirtz, Hanno

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/038823

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2011/292072 A1 (FISHER DANYEL [US] ET AL) 1 December 2011 (2011-12-01) abstract figures 1, 4 paragraph [0004] - paragraph [0005] paragraph [0014] - paragraph [0033] paragraph [0040] - paragraph [0055] -----	1-15
X	US 2006/218563 A1 (GRINSTEIN GEORGES [US] ET AL) 28 September 2006 (2006-09-28) figures 1, 5 paragraph [0006] - paragraph [0041] paragraph [0093] - paragraph [0109] -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2019/038823

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011292072	A1	01-12-2011	NONE

US 2006218563	A1	28-09-2006	US 2006218563 A1 28-09-2006
		US 2010268691 A1	21-10-2010
