

US PATENT & TRADEMARK OFFICE

PATENT APPLICATION FULL TEXT AND IMAGE DATABASE



(1 of 1)

United States Patent Application**20070249368****Kind Code****A1****Bailly; Olivier ; et al.****October 25, 2007**

Shared Geo-Located Objects

Abstract

A method of reviewing geo-coded information at a geographic information system is disclosed. The method may include receiving information identifying a geo-located object from a submitter of the object, receiving information identifying the location of the geo-located object from the submitter, and receiving a request from a user of the geographic information system for one or more geo-located objects that includes the submitted geo-located object, and transmitting information for permitting display of the geo-located object to the user of the geographic information system.

Inventors: **Bailly; Olivier; (Oakland, CA) ; Bailly; Francois; (Santa Clara, CA) ; Smith; Adam; (San Francisco, CA) ; Limber; Mark; (Boulder, CO) ; Springer; Michael W.; (Aurora, CO)**

Correspondence Name and Address: **FISH & RICHARDSON P.C.
PO BOX 1022
MINNEAPOLIS
MN
55440-1022
US**

Assignee Name and Address: **GOOGLE INC.
1600 Amphitheatre Parkway
Mountain View
CA
94043**

Serial No.: **740222**Series Code: **11**Filed: **April 25, 2007**

U.S. Current Class: 455/457
U.S. Class at Publication: 455/457
Intern'l Class: H04Q 7/20 20060101 H04Q007/20

Claims

1. A method, comprising: receiving, through an internet-accessible interface, information identifying a geo-located model from a submitter of the object; receiving information identifying the location of the geo-located model from the submitter; storing the geo-located model in a model database with a plurality of other geo-located models; and forming an association between image information for a geography with the geo-located model.
2. The method of claim 1, further comprising receiving a request including information identifying a geographical location, and causing image information for the geographical location and information defining a model associated with the geographical location to be delivered to a client device submitting the request.
3. The method of claim 1, wherein the information defining a model associated with the geographical location is transmitted in an XML format.
4. The method of claim 1, wherein the information identifying the location of geo-located model includes a placemark that identifies the location of the model and a link directed to the model, wherein selection of the link causes the model to be downloaded to a requesting client device.
5. The method of claim 4, further comprising storing and transmitting an image of the model, wherein the image is formatted so that selection of the image by a user causes the model to be displayed on a client device associated with the user.
6. The method of claim 1, further comprising associating advertising with the geo-located model for display on the geo-located model.
7. A method, comprising: receiving a selection of a geographical location and an identification of a geo-located model; transmitting to a central server information identifying the geographical location and information identifying a location of information defining the geo-located model; and providing information defining the geo-located model to a model server that makes the information defining the geo-located model publicly available.
8. The method of claim 7, further comprising generating one or more images of the geo-located model, and providing the one or more images for display in response to search requests for model information.
9. The method of claim 7, further comprising generating tags descriptive of the geo-located model, and providing the tags as keywords for public requests for model information.
10. The method of claim 7, wherein the geographical location represents a portion of a whole earth representation.
12. The method of claim 9, wherein the model server includes an interface to receive public submissions associated with geo-correlated models and make the submitted models publicly available.

13. A computer-implemented system, comprising: an interface generating requests for, and receiving, information identifying a geographic location and a geo-correlated model; a model server associated with the interface that stores information associated with the geo-correlated model for display with the geographical areas; and a correlation structure associating the geographical areas with the one or more models, and configured to receive geo-coded submissions from users submitting models to the system.

14. The system of claim 13, wherein the geo-coded model includes a three-dimensional model of a man-made structure.

Description

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority under 35 USC .sectn.119(e) to U.S. Patent Application Ser. No. 60/794,875, filed on Apr. 25, 2006, the entire contents of which are hereby incorporated by reference.

TECHNICAL FIELD

[0002] This invention relates to data storage in distributed systems, and more particularly to databases for geographic information systems.

BACKGROUND

[0003] Geographic information systems (GIS) permit for the archiving, retrieving, and manipulation of data that has been stored and indexed according to the geographic coordinates of its elements.

[0004] GIS systems can provide very helpful information to users, such as by presenting information graphically--see Google Maps and Google Earth as two prime examples of such helpful systems. GIS systems often are navigated by entering particular longitude/latitude coordinates, or by entering addresses, such as street addresses, which may be resolved to a particular coordinate before querying a database for geo-specific information such as maps or satellite images. The information returned may be information that is provided by "official" sources, such as government or industry imagery or demographic data, or information such as the locations of businesses and residences.

SUMMARY

[0005] This document describes systems and methods for storing and accessing geo-located content, such as 3D content to be displayed on a model of the earth. In general, users of a system may identify objects to be placed in a geography to a system. These objects may include, for example, 3D models of buildings and other structures--whether current, planned, or fictional. The system may store the objects or may simply store information pointing to the objects. The same or other users may then navigate to certain locations in the geography, or may select particular objects or classes of objects. Information pertaining to the objects may then be shown in a display that also shows the geography, and the information about the objects may, for example, be overlaid over the geography in the locations of the objects.

[0006] In response to a selection by the viewing user, such as selecting a labeled hyperlink, the system may return additional information about the object. For example, the system may return an image of the

object for the user to preview. The system may also present non-graphical information about the object, such as by opening a separate browser window to display data, e.g., data about earthquakes in certain areas of a geography. Alternatively, the system may point the user's application at the object itself. Where the object is a 3D model, information encoding the object may then be passed to the user's application, which may then render the object on the geography, for interaction with the user.

[0007] In one implementation, a method of reviewing geo-coded information at a geographic information system is disclosed. The method may receiving information identifying a geo-located object from a submitter of the object, receiving information identifying the location of the geo-located object from the submitter, receiving a request from a user of the geographic information system for one or more geo-located objects that includes the submitted geo-located object, and transmitting information for permitting display of the geo-located object to the user of the geographic information system.

[0008] The described systems and methods may provide for one or more of the following features and/or advantages. Users of a system may benefit by being able to locate, preview, display, and interact easily with objects in a geography such as that displayed by Google Earth. Producers of objects may easily associate their objects with a geography and thereby publish their works so that others may access them. For example, architects may construct models of buildings and make them available to clients through an application such as Google Earth. Geographic information providers may benefit by drawing additional users to their applications, and may thereby increase the revenue they can generate with such applications.

[0009] The details of one or more embodiments are set forth in the accompanying drawings and the description below. Other features, objects, and advantages will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

[0010] FIG. 1 is a schematic diagram of a system for providing users with objects in a geographic information system.

[0011] FIG. 2 is a swim-lane diagram of actions that may be taken to provide a user with information about objects associate with a geography.

[0012] FIG. 3 shows relationships that may be formed between versions of a model for a geographic information system.

[0013] FIG. 4 is a flowchart for uploading objects to a geo-coded database.

[0014] FIG. 5 is a schematic diagram showing components in a computer system suitable to be used with the systems and methods described in this document.

[0015] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0016] FIG. 1 is a schematic diagram of a system 100 for providing users with objects in a geographic information system. In general, the system 100 uses a geography server 114 to generate views of a geographic area on an application running on terminal 110--such as a web browser (with or without a plug-in directed to the geographic display) or other application (such as Google Earth). The system 100 may also in cooperation use an object server 116 to deliver information concerning objects associated

with the geography. Such objects may include 3D models of building, bridges, or other structures, and may also include geographic data such as earthquake data, weather data, or the like. The coordinated use of the servers 114, 116 may allow objects to be layered over the top of geography so that users may more readily review such objects in a natural manner, and may also interact with the objects in helpful manners.

[0017] The various components of the system 100 may communicate through network 112, which may include one or more local area networks (LANs), wide area networks (WANs), the internet, or other appropriate networking structures. The particular arrangement of network 112 is not critical.

[0018] Geography server 114 may store a geographic model in geography database 115, and may use a mapping module 122 to receive and respond to user requests through an interface 118. The interface 118 may be, or be part of, a web server or other similar structure or structures. The geography database 115 may include one or more database that contain information showing geographic information, such as a map, satellite images, or visual terrain (e.g., satellite images applied to topography).

[0019] The mapping module 122 may receive requests from users, such as a user at terminal 110 through interface 118, and may determine what information in geography database 115 the user needs. For example, the interface 118 may pass certain location information, such as an identifier of a latitude and longitude at which the user has located his or her application, along with information that permits the determination of a bounding box for the area viewed by the user. The mapping module 122 may then query geography database 115 to obtain data for the area inside the bounding box or other defined area. The mapping module 122 may then pass the information to terminal 110, such as in the form of tiled geographic information around the area the user is viewing, or in another appropriate format.

[0020] The object server 116 may contain an object database 117 that stores information about various objects that may be associated with particular areas in a geography. The object database 117 may store, for example, a geographic location identifier for each object so that it may be located within, and overlaid on, a representation of the geography. The database 117 may also store information about the object that may be displayed in a geographic representation on terminal 110, including in the form of hyperlinks to additional information. For example, the database 117 may include a label, such as "Eiffel Tower models" that may be displayed over a map of Paris. When a user selects the label, which may be in the form of a hyperlink, the user may be presented with one or more images showing various 3D models of the Eiffel Tower, such as in a pop-up window. (The models may be ordered in various manners, such as by placing most popular models near the top or placing models from certain "trusted" object builders at the top.) The user may select one of the models, such as by selecting one of the images, and the model may be added to the geography displayed on terminal 110.

[0021] Object selection engine 124 may operate through interface 120 to identify objects to be displayed on terminal 110. Interface 120 may again be a web server or servers, or may be part of a web server or other appropriate structure. Object server 116 may receive the same commands as does geography server 114, and may, in parallel with geography server 114, obtain information about objects to be sent to terminal 110. The object selection engine 124 may, for example, receive information about a location (including information about a boundary for the location, or information that permits the boundary to be determined), and may query object database 117 to find appropriate objects in the area. The objects may be classified into particular groups, such as types of buildings, weather information, etc., and subsets of database 117 may be searched based on particular groups selected by a user of terminal 110.

[0022] Object information obtained from object database 117 may be provided in a variety of forms, including via extensible mark-up language (XML) or an XML extension known as Keyhole Mark-up Language (KML). KML is an XML grammar and file format for modeling and storing geographic

features for display in the Google Earth Client application. Particular definitions for version 2.0 of KML are shown in the attached Appendix A.

[0023] Labeled arrows that are layered over the components of system 100 show an exemplary flow of information in the system 100. Arrow 1 from rack server 128 shows the submission of information about a particular object to system 100. Rack server 128 represents, for example, a commercial object generator, such as an architectural firm or other such firm. Such an object generator might also produce objects, such as 3D building models, so as to be able to place advertising on the buildings and thereby generate revenue from users who view such objects.

[0024] Server 130 may also submit objects to system 100. Server 130 represents, for example, smaller object providers. As one example, a user of a computer may generate an object using an application such as SketchUp, AutoCAD, or 3DStudio, and may submit the object to system 100 for display as part of a geographic information system. In one implementation, the user may visit a web site associated with system 100, and provide information about the object. For example, the user of server 130 may provide a URL where the object may be accessed, such as when other users later choose to incorporate the object into geographies they are viewing. In addition, the user may provide additional information, such as an exemplary image of the object, and other information about the object (e.g., labels for hyperlinks associated with the object). Rack server 128 may provide similar information when submitting an object, and may do so via an interactive web form, or via another application programming interface.

[0025] Information about objects submitted to system 100 may be received by object server 116, which may parse such submissions and store the relevant information in object database 117. The information may include a geographic identifier for the object, so that the object may be readily identified later when a user at client 110 is viewing geography, and requests to see objects associated with that geography.

[0026] The lettered flow arrows show actions that may occur when a user at terminal 110 begins looking at information in a geographic information system. At arrow A, the user navigates to a particular area of geography, so that a request for visual information about that area is sent through network, as further indicated by Arrow B, to geography server 114. The message may also be routed to object server 116, if the user has selected to have displayed any categories of objects associated with the geography. Geography server 114 determines what information is needed to provide interaction for the user with the particular area, and returns that information via Arrow C and Arrow D. The terminal 110 may then begin constructing a view of the area as the information is downloaded.

[0027] At Arrow E, the user may select to view additional information about an object displayed in the area of geography. For example, the user may see a link representing a model of a monument (or group of monuments), and may elect to see the model or models. As one example, a user at rack server 128 may have modeled the monuments on and around the National Mall in Washington D.C., and the terminal 110 may show a link to such models. The user at terminal 110 may select a link displayed over a geographic representation of the Mall (e.g., a map or satellite photo) to obtain more information about the models.

[0028] Arrow F shows the user's request being received by object server 116. The request may be formed, for example, in the form of an HTTP request or an XML message. An identifier for the object may be parsed out of the request by server 116, for example, and submitted to object database 117 as a query. The database 117 may then return information, such as an image of the models, which may be returned to terminal 110 via Arrow G and Arrow H.

[0029] If the user of terminal 110 likes the models shown in the image, he or she may select to receive the full models. That selection may occur, for example, by clicking on the image. Such an action may

cause a message to be delivered to rack server 128 via Arrow I and Arrow J, and rack server may respond by sending information needed to display the model or models via Arrow K and Arrow L.

[0030] FIG. 2 is a swim-lane diagram of actions that may be taken to provide a user with information about objects associate with a geography. At act 200, a user of a client device navigates to a geographic location, such as by entering a lat/long combination or an address, or by moving a graphical representation of the geographic area into view on a display. Messages are sent to a geography server and an object server, and may define the current viewport for the application running on the user's terminal. The message may take the form, for example, of a query parameter that identifies a viewport showing a particular geographic area. In one example, if a user is viewing the city of San Francisco, the message may take the form of: [0031] Bbox=37.9%2C37.7%2C-122.2%2C-122.4

[0032] The geography server returns a representation of the location, such as by obtaining satellite or map tiles for the area and for a geography a set distance around the area (act 202). The object server uses geographic information received from the client to look-up and return objects that may be in the requested area/location (act 204).

[0033] The information returned by the object server may be, as noted above, textual or image information, along with hyperlink information relating to architectural or other objects in the geographic area. The information may be transmitted, as indicated above, in a KML document, which may contain, for each object or model: [0034] A placemark that pin-points the location of the model on the earth; [0035] A link to a preview image of a thumbnail of the model; and [0036] Links to allow a user to download the model, such as in applications like Google Earth and Google SketchUp or other display and editing applications.

[0037] At act 206, the client displays the geographic information at the particular locations, along with object indicators, such as hyperlinks. The indicators may be displayed, for example, overlaid on the geographic information, or to the side of the geographic information, such as in a list of objects relevant to the display. A user at the client may then select an object indicator 208, such as an indicator for a preview image of an object (act 208). This selection may cause the object server (or another appropriate structure) to deliver a preview image to the client (act 210).

[0038] The user at the client may then select to view more complete information about an object (act 212), such as information needed to define a full 3D model. Such selection may occur, for example, by selecting an image of the model having a hyperlink associated with it. The link may send a message to the object server, or to another server that requests information about the model or models. The server may then retrieve information about the model or models and transmit it back to the client (act 214).

[0039] Upon receiving the information, the client may use the information to locate and render the object for visual review by a user (act 216). For example, the client may render the object as a 3D item on top of a representation of the earth.

[0040] FIG. 3 shows relationships that may be formed between versions of a model for a geographic information system. In particular, where models are available for access as described above, users may search for and obtain models of interest and download them in a manner that allows for further editing of the models. In the figure, the exemplary model is a simple model of a farmhouse.

[0041] Modifications of the model by various users are shown hierarchically in the figure. The so-called "root" model is Obj1. It is a very basic massing model of single-room house, containing no details. At Obj1a, another user has accessed the model and added details (e.g., windows, a door, and a chimney) but has left the massing alone. At Obj1b, another user has accessed the model, added a new mass, and added

two doors. A further user has accessed Obj1b, reversed the massing, and added windows, to create Obj1b1. Obj1a has been amended at Obj1a1 (which has further been altered at Obj1a1a) and Obj1a2.

[0042] These representations of various revisions on a model show changes that various users may make when basic models are posted publicly and made available for further alternation. When a user accesses or downloads a model, a system may create a flag indicating the model and the user. The user may then return the model after making additions or other alterations to it, and the flag may be used to connect the newly submitted model to the base model that the user originally accessed. In this manner, logical connections may be formed between various versions of a model so as to permit ready access to a particular model or a particular version of the model.

[0043] The models and the relationships between and among the models may be displayed graphically to a user in a manner similar to that shown in FIG. 3. As such, when a user navigates to a certain geography, the user may be presented with a number of different objects in the area. Objects that are placed in the same point or in overlapping areas may be considered, in certain situations, to be models of the same object. For example, various users may have created their own renditions of the Statue of Liberty. A user may select the particular object and be presented with images of each of the models, which may be "root" level versions of a model. Where multiple versions of a model exist, those versions may be shown in a manner like that in FIG. 3 once a user selects that particular model to retrieve. The user may then be given a choice to select a particular version of the model. For example, perhaps the root version is overly simplified for the user's tastes. Or perhaps the root version shows an actual structure, whereas the user wishes to view a fanciful extension of the model, such as a Statue of Liberty having the form of Larry King.

[0044] FIG. 4 is a flowchart for uploading objects to a geo-coded database. At act 402, a user of a client device (or an application running on the client device) determines that an object may be uploaded to a geographic information system. For example, the user may have constructed a model of a building in a program such as SketchUp, and may determine that they would like to have the model be accessible to others in part of a map or other geographic representation.

[0045] At act 404, a server determines whether the user is currently logged into a system. For example, a user may be permitted a single log in during one computing session, and may use that log in to access a number of applications provided in combination with the server, such as e-mail, text messaging, and other applications. As used here, the server may be a single machine or multiple computing devices organized in a manner so as to carry out a number of functions associated with the system.

[0046] If the user is not determined to be currently logged in, such as by checking an identifier sent with the upload request by the user, the system may check to determine whether the user has an active account with the system (act 406). If not, the user may go through a process for activating a new account (act 408) and it may then be determined whether the user has agreed to terms of service for the system and/or the particular application (act 412). If the user has a valid account (act 406), then it is determined whether the user has a valid log in, e.g., the user's credentials may be checked by prompting the user for a UserID and a password (act 410).

[0047] When a user has been authenticated by whatever means, the system may check to see if the user has yet agreed to the terms of service for uploading objects. If the user has not yet accepted the terms of service, he or she may be presented with a terms of service form (act 414). For example, the user may be reminded that any content they upload can be viewed, used, and modified by other users without any intellectual property violations by the system or by others. Also, the user may be provided with a notice that the uploaded content will be subject to a Creative Commons license, for use by other users. The form may require acceptance by the user before the user may go any farther, and may additionally

require an active acceptance, such as by blanking an acceptance box and requiring the user to check the box before continuing. Agreement to the terms of service may be made mandatory for uploading models, but not for viewing models.

[0048] At act 416, the user is asked to add meta data relating to their object. For example, a user may provide a URL that points to the model on their computing device so that, for example, the system may redirect other users there or may obtain the object from that location for storage on the system's servers. The user may also provide an image of the object for later display to other users. In addition, the user may identify other users or classes of users who may have access to the model, and may define levels of permissible access. For example, a user may identify an object as "read only" or "view only" so that others may not edit it. The system may place restrictions on the ability of users to place such restrictions, however, as narrow sharing of designs does not encourage later user interaction.

[0049] In addition, a user may identify a password and private identifier for an object. In this manner, the user may make the object available only to a select group of other users. For example, an architect may want to prepare a model for review only by a client, but place the model in its proper location within a city. The architect may also accompany the model with a number of mere massing, gray models of other buildings around the design, so that the proposed design is a centerpiece, but can be reviewed in its proper context by clients. However, the architect may not want the public to see the model yet, and may thus password-protect access to the model. Either through the server, or otherwise, the architect may identify e-mail addresses of the users who the architect would like to see the model, and may have invitations delivered to those users. The invitations may include a URL that when selected by the users, causes an application to launch or download for displaying the model in its proper geography. Passwords may be sent separately to the users. To encourage more public use of the system, the use of password protected models may be limited by charging for such use, and making public use free. Alternatively, users of password-protected models may be required to review additional paid promotional material over-and-above what other users would have to view.

[0050] At act 418, the server confirms that a model or other object has been properly uploaded or linked to the system. Although such feedback may be given immediately, the availability of the object in the system may occur at a later time. In contrast, if the user cancels out of the process, the server or client may cause a user's window to be closed without finishing the transaction.

[0051] Uploads may be sent to a Java (GSE) frontend, which may do simple validation on KML coming from model-makers to check that it is syntactically correct and extract its bounding box. Assuming that the file is valid, the frontend will write the KML file, SKP file, bounding box, and any associated keywords to a single row in a bigtable, which is a system for managing large amounts of data distributed across many machines.

[0052] A C++ backend server may periodically scan the bigtable and build up an in-memory index that can be used for keyword and bounding-box searches. The frontend can send RPC search queries to the backend, and the backend may return search results that the frontend will then display. This backend may also serve VBR queries, and return KML placemarks for all models inside a query bounding box. When a user requests a model by clicking on the URL in the placemark, the request may be sent to the frontend server, so that all bandwidth-intensive requests go to the same place.

[0053] Various messages may be shared between the front-end and the back-end, including messages relating to metadata associated with each entry (e.g., title, description, author userId, etc.), the popularity of an entry (e.g., tracks number of clicks and download over the past seven days), parameters sent by the frontend for querying a model, messages that associate a metadata with its row name in the bigtable as well as its popularity, a message returned by the backend for queries by lat/long that indicates that many

objects are available under an area which would be available at a higher zoom level, messages sent by the front-end to notify the back-end of new/modified data, messages sent by the front-end to update popularity of a list of items, and response messages for queries or index updates.

[0054] The front-end may be provided with a number of Java servlets, such as servlets permitting keyword-based search; a servlet to return details on a particular model; a servlet for uploading, such as to validate the content, and if successful, write it to bigtable; a page where the user fills in the details on the model being submitted; a download servlet that retrieves a KML file from the bigtable and returns it; a servlet to delete a model; a servlet to serve cached model thumbnails; and a servlet to manage a terms of service page. The upload functionality may also be part of a model design program, such as SketchUp, itself, and may be provided as a menu option in such an application.

[0055] Uploaded models may be written to the bigtable after verification and bounds extraction. The number of polygons in each model may also be counted and stored as meta data. For each new model, the front-end may compute a unique rowname (see below) and create the metadata to store in the bigtable's first column. The other two columns may contain the KMZ and SKP files. The row name for each new model may be computed as follows: [0056] a random 8-byte number XOR'ed with the current time in milliseconds, represented on 8 bytes as well (using `System.currentTimeMillis()`) [0057] The hash value (8 bytes) of the string representation of the model's metadata, concatenated with the requestor's IP address. [0058] Both 8 byte numbers written as a hex string using `Integer.toHexString()` and concatenated to produce a 32-character identifier.

[0059] Note that the row name may contain no information related to the geographic position of the model. Location of the model can be deduced from its computed bounding box, which is present in the model metadata. Once the row name is computed, the model metadata, KMZ file and SKP files may be written to the bigtable (which may be a very large database table) in the appropriate columns (metatada, kmz and skp).

[0060] From in an application such as SketchUp, a user may first choose to upload a model, and the application may then launch a web dialog with a simple request URI (no parameters or post data). Once the user is authenticated, a submit page may be displayed to permit entry of a title, description, comma-separated tags, and browsing for an optional user-created image of the model.

[0061] When the user submits the necessary information, they may then click to submit the model. The submit button may point to URL ``skp:upload@/submit``, where ``submit`` is an actual upload URI. The application may intercept the submission, package up the files (e.g., into a zip containing .skp, .kml, thumbnail.jpg, optional user-created image) into a single POST, and submit it to the upload URI in the same browser session. The request URI may take the form of: [0062] `/upload?title=[...]&description=[...]&tags=[...]&bbox=[W,S,E,N,low,hi gh]&polygons=[...]`

[0063] The server may then respond with a success or a failure page, depending on what happened during upload.

[0064] Updates to An Existing Model

[0065] Optionally, clients (such as Sketchup) can identify the model being uploaded with its id (bigtable rowname) in order to submit updates to existing models. Assuming the user trying to update a model is also the model's author, the upload info page will be prepopulated with the metadata for the existing model. An application may extract the ID of a model by querying the value of a hidden input field in the web page presented when the model is uploaded, and store that id in the model itself after a successful upload.

[0066] The web interface also has links to allow a model's author to edit the model's metadata from the web interface (useful to correct errors like typos, or entering better tags to categorize models). Models may be made available to users looking at a certain geography, or may alternatively be presented as a result to a search, such as a text search. The search may be a traditional keyword-based search, or may be user-based searches looking for models by a particular user. Searching may also look for tags that identify features relating to models. The front-end may also supports searching for the most recently uploaded and most popular models.

[0067] Users may be allowed to flag items as bad or offensive by clicking on a "Report bad item" link located next to each model in a results and details page. That link will redirect (using a URL that encodes information that allows identification of the model) the user to a page that permits entry of particular data by the user for identifying why the model may be offensive.

[0068] A model's author may delete their own models from the system. So-called "deleted" models may initially be hidden but not immediately deleted from the bigtable. In this manner, users may be provided with a link to their hidden models, and be given the opportunity to reactivate them.

[0069] The backend may periodically scan the keywords in the bigtable and build an in-memory index to answer searches. Ranking may use popularity (user downloads). The front-end can keep popularity counts updated in the bigtable, and the indexer can periodically rescan the table to update its knowledge of the counts.

[0070] Shape Indexing: Users may also be permitted to search for particular shapes. For example, a model may be viewed as consisting of a number of "primitive" shapes such as spheres, 3D rectangles, toroids, and the like, and models may be composed of Boolean combinations of these primitives. A model may be associated with one or more shapes, such as by assigning a score for particular shapes to the model, by, for example, how many times a particular shape occurs in a model, and how prominent (e.g., in terms of size) the shapes are when they appear. Also, shapes may be classified by the number of parallel surfaces they have, with spare shapes characterized by straight lines/surfaces, and parallel lines/surfaces--and ornamented designs characterized by more complex surfaces.

[0071] Each row in the bigtable may contain the data for one model. Row names may be hashes of various parameters of the model, and need not have inherent meaning. Metadata may be stored separately from other model data so that metadata scans avoid loading the large binary data.

[0072] For the object submission, a client such as the SketchUp Client may display an embedded web browser, and points it to a URL for an Upload Metadata page on a server hosting a 3D object database. On this page, a user may enter metadata to be attached to the uploaded model. Such metadata fields may include: [0073] title [0074] description (free-form text) [0075] URL (user-supplied URL where more details can be found about the uploaded model) [0076] tags (list of words/phrases that can be used to categorize the model)

[0077] Additionally, if the model is geo-located, then the request URL may contain a request parameter that signifies this. In this case, the 3D object server may add the following fields to the upload metadata page: [0078] physical address (actual location of the the physical structure the model represents) [0079] model type (e.g., "real," "fantasy," "planned")

[0080] In one implementation, the upload metadata page may contains an Upload button that points to a URL of the form "skp:upload@<UPLOAD-URL>," where <UPLOAD-URL> is the actual URL to which the model should be posted. When the user clicks the Upload button, the client may intercepts the

"skp:" URL and (1) if the model is geo-located, export it to KMZ (the format readable by Google Earth); (2) generate a preview image for the file, in JPEG format; (3) create a Zip archive containing the current SketchUp-format model, the preview image, and the exported KMZ (the latter only for geo-located models); (4) send the Zip archive to the <UPLOAD-URL> via HTTP POST, with the names/values of the metadata fields added to the URL as query parameters; and (5) the 3D database server may receive the POST, extract the Zip archive, and store all of the submitted files and metadata.

[0081] In another implementation, the upload metadata page may allow a user to submit all of the fields mentioned above by pressing an Upload button. If the values are invalid in any way, the upload metadata page may be presented again with errors highlighted. If the metadata values are accepted, the following actions may occur: [0082] The 3D database server stores the metadata values and generates a Unique ID that can be used to look them up. [0083] The 3D database server sends a browser redirect to the client. The redirect URL is of the form "skp:upload@<UPLOAD-URL>," where <UPLOAD-URL> is the actual URL to which the model should be posted. The URL contains the metadata's Unique ID as a query parameter. The Client intercepts the "skp:" URL and does the following: (1) if the model is geo-located, exports it to KMZ (the format readable by Google Earth); (2) generates a preview image for the file, in JPEG format; (3) creates a Zip archive containing the current SketchUp-format model, the preview image, and the exported KMZ (the latter only for geo-located models); (4) sends the Zip archive to the <UPLOAD-URL> via HTTP POST; and (5) the 3D database server receives the POST, extracts the Zip archive, and stores all the extracted files, associated with the originally-submitted metadata using the Unique ID in the <UPLOAD-URL>.

[0084] Models can be located and obtained by the 3D database server, using the Get Models feature in a SketchUp Client. The SketchUp Client may display an embedded web browser that points to a Search URL on a 3D database server. Models stored on the server may be presented to the user in a standard search/results/details manner.

[0085] Each model detail page may also contain a Download to SketchUp button. The URL for this button may be of the form "skp:download@<DOWNLOAD-URL>," where <DOWNLOAD-URL> is the actual URL from which the SketchUp model can be obtained. When the button is clicked, the following may occur: [0086] The client intercepts the "skp:" URL and displays a dialog asking whether or not to download directly into the current document. If the answer is No, the embedded web browser itself may be redirected to the <DOWNLOAD-URL>, and a standard browser download may take place. If the answer is Yes, then the client may make an HTTP request to download the model. It may store the downloaded model in memory and load it directly into the current design document. [0087] If the model was downloaded directly into the current document, there may be several behaviors that depend on whether the model is geo-located: [0088] If the model is not geo-located, then it may be pushed onto a tool that allows the user to choose the 3D location of the model graphically. [0089] If the downloaded model is geospecific, then: [0090] If the_current_model is geospecific, then the user is prompted to either put the downloaded model in its "natural" place relative to the current one, or to place it manually using the Place Component tool. [0091] Otherwise, then the downloaded model's geo-location information is transferred to the current model (making it geo-located).

[0092] FIG. 5 is a schematic diagram showing components in a computer system suitable to be used with the systems and methods described in this document. The system 500 includes a processor 510, a memory 520, a storage device 530, and an input/output device 540. Each of the components 510, 520, 530, and 540 are interconnected using a system bus 550. The processor 510 is capable of processing instructions for execution within the system 2300. In one implementation, the processor 510 is a single-threaded processor. In another implementation, the processor 510 is a multi-threaded processor. The processor 510 is capable of processing instructions stored in the memory 520 or on the storage device 530 to display graphical information for a user interface on the input/output device 540.

[0093] The memory 520 stores information within the system 500. In one implementation, the memory 520 is a computer-readable medium. In one implementation, the memory 520 is a volatile memory unit. In another implementation, the memory 520 is a non-volatile memory unit.

[0094] The storage device 530 is capable of providing mass storage for the system 500. In one implementation, the storage device 530 is a computer-readable medium. In various different implementations, the storage device 530 may be a floppy disk device, a hard disk device, an optical disk device, or a tape device.

[0095] The input/output device 540 provides input/output operations for the system 500. In one implementation, the input/output device 540 includes a keyboard and/or pointing device. In another implementation, the input/output device 540 includes a display unit for displaying graphical user interfaces.

[0096] The features described can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations of them. The apparatus can be implemented in a computer program product tangibly embodied in an information carrier, e.g., in a machine-readable storage device or in a propagated signal, for execution by a programmable processor; and method steps can be performed by a programmable processor executing a program of instructions to perform functions of the described implementations by operating on input data and generating output. The described features can be implemented advantageously in one or more computer programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. A computer program is a set of instructions that can be used, directly or indirectly, in a computer to perform a certain activity or bring about a certain result. A computer program can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment.

[0097] Suitable processors for the execution of a program of instructions include, by way of example, both general and special purpose microprocessors, and the sole processor or one of multiple processors of any kind of computer. Generally, a processor will receive instructions and data from a read-only memory or a random access memory or both. The essential elements of a computer are a processor for executing instructions and one or more memories for storing instructions and data. Generally, a computer will also include, or be operatively coupled to communicate with, one or more mass storage devices for storing data files; such devices include magnetic disks, such as internal hard disks and removable disks; magneto-optical disks; and optical disks. Storage devices suitable for tangibly embodying computer program instructions and data include all forms of non-volatile memory, including by way of example semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, ASICs (application-specific integrated circuits).

[0098] To provide for interaction with a user, the features can be implemented on a computer having a display device such as a CRT (cathode ray tube) or LCD (liquid crystal display) monitor for displaying information to the user and a keyboard and a pointing device such as a mouse or a trackball by which the user can provide input to the computer.

[0099] The features can be implemented in a computer system that includes a back-end component, such as a data server, or that includes a middleware component, such as an application server or an Internet

server, or that includes a front-end component, such as a client computer having a graphical user interface or an Internet browser, or any combination of them. The components of the system can be connected by any form or medium of digital data communication such as a communication network. Examples of communication networks include, e.g., a LAN, a WAN, and the computers and networks forming the Internet.

[0100] The computer system can include clients and servers. A client and server are generally remote from each other and typically interact through a network, such as the described one. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0101] A number of embodiments have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of what is described. Accordingly, other embodiments are within the scope of the following claims.

* * * * *

