



- (51) International Patent Classification:  
*G06F 17/40* (2006.01)    *G06F 15/16* (2006.01)
- (21) International Application Number:  
PCT/US2012/022577
- (22) International Filing Date:  
25 January 2012 (25.01.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
13/014,689    26 January 2011 (26.01.2011)    US
- (71) Applicant (for all designated States except US):  
**GOOGLE INC.** [US/US]; 1600 Amphitheatre Parkway,  
Mountain View, California 94303 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **APPLETON, Benjamin C.** [AU/AU]; 41 Spencer Street, Summer Hill, New  
South Wales 2130 (AU).

- (74) Agents: **SHEPHERD, Michael P.** et al.; Fish & Richardson P.C., P.O. Box 1022, 3300 Rbc Plaza, Minneapolis, MN 55440-1022 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, 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, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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,

[Continued on next page]

(54) Title: CACHING RESOURCES




FIG. 2

(57) Abstract: Methods, systems, and apparatus, including computer programs encoded on a computer storage medium, for caching tiles in an interactive mapping system. A request is received from a first client, the request being for a root node of a directory structure in which resource locators for nodes are generated based on a hash of resource locators of respective descendant nodes, and wherein leaf nodes are associated with corresponding resource locators for tiles in an interactive mapping system. A first configuration of tiles is served to the first client as cacheable resources. A new node is added to the directory structure corresponding to a new version of a tile. A second configuration of tiles is served to a second client, while continuing to serve the first configuration of tiles to clients that requested the root node before the resource locators for ancestor nodes of the new node were added.

WO 2012/103237 A1

LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, **Published:**  
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, — *with international search report (Art. 21(3))*  
GW, ML, MR, NE, SN, TD, TG).

## CACHING RESOURCES

### BACKGROUND

This specification relates generally to caching resources in a network.

5 In a typical interactive mapping system, tiles of map data are served over a network through multiple intermediary servers. Tiles may be cached on each of the intermediary servers. For a tile that has been cached, subsequent requests for that tile result in a local copy of the cached tile being served to a client.

### SUMMARY

10 This specification describes technologies relating to caching resources in an interactive mapping system.

In general, a directory structure is created by the interactive mapping system to control caching of tiles. Clients traverse the directory structure by making subsequent requests for resource locators at each level of the directory structure to ultimately obtain tiles corresponding to a set of map coordinates. When a new tile is added to the system, 15 only ancestor nodes of the new tile are updated. Thus, while new tiles and their respective ancestor nodes are uncached, unchanged tiles remain cached.

In general, one innovative aspect of the subject matter described in this specification can be embodied in methods that include the actions of receiving, from a first client, a request for a root node of a directory structure in which resource locators for 20 nodes are generated based on a hash of resource locators of respective one or more descendant nodes, wherein requests for parent nodes generate responses containing resource locators of respective one or more descendant nodes, and wherein leaf nodes are associated with corresponding resource locators for tiles in an interactive mapping system; serving, to the first client, a first configuration of tiles, wherein each intermediate 25 node and each tile is served as a cacheable resource; receiving an indication of a new node added to the directory structure, the new node corresponding to a new version of a tile; adding to the directory structure one or more new ancestor nodes of the new node; receiving, from a second client, a request for the root node; and serving, to the second client, a second configuration of tiles including the new node, while continuing to serve 30 the first configuration of tiles to one or more other clients that requested the root node before the one or more new ancestor nodes of the new node were added. Other embodiments of this aspect include corresponding systems, apparatus, and computer

programs, configured to perform the actions of the methods, encoded on computer storage devices.

5 These and other embodiments can each optionally include one or more of the following features. Adding one or more new ancestor nodes of the new node includes  
switching from serving the first configuration of tiles to serving the second configuration  
10 of tiles after a new root node is added. Switching to serving the second configuration of tiles includes swapping a root node indicator from the root node to the new root node. A response containing resource locators of one or more descendant nodes of the new root node is served in response to receiving, from a second client, a request for the root node.  
15 The resource locators for tiles include a version. The second configuration of tiles is served with an indication that one or more tiles are cacheable. Updating resource locators for each of one or more ancestor nodes of the new node comprises adding to the directory structure the one or more ancestor nodes for the new node. The resource locators for tiles include map coordinates. The resource locators for tiles include a hash of tile data.  
20 Resource locators for the root node and each parent node are generated by a hash function. Resource locators for the root node and each parent node are generated by a hash of a concatenation of resource locators of one or more respective descendant nodes. A separate directory structure is generated for each of one or more zoom levels in the interactive mapping system. The directory structure is a quadtree or a B-tree. The  
method of claim 1, where the indication that the tile is cacheable comprises an HTTP header.

Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. The directory  
25 structure can reduce latency for client devices and reduce load on servers by increasing a cache hit rate. New tiles can be served without simultaneously invalidating all caches for all users. Tiles need not be specified by predictable uniform resource locators (“URLs”), and can instead be specified by arbitrary URLs, such as those based on tile contents. URLs based on tile contents result in identical URLs being assigned to visually identical tiles, further increasing cache hit rates throughout the system.

30 The details of one or more embodiments of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

## BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram of a graphical user interface of an example interactive mapping system.

FIG. 2 is a diagram of an example network environment for serving map data.

5 FIG. 3 is a diagram of an example directory structure for caching tiles.

FIG. 4 is a sequence diagram of an example client interaction with the interactive mapping system.

FIG. 5 is a diagram of an example update of the directory structure for caching tiles.

10 FIG. 6 is a sequence diagram of an example client interaction with the interactive mapping system after a tile is updated.

Like reference numbers and designations in the various drawings indicate like elements.

## DETAILED DESCRIPTION

15 Interactive mapping systems provide access to vast amounts of map data, particularly when provided in a networked environment, e.g., the Internet. The interactive mapping systems can store the map data in a distributed storage environment and serve the map data to client devices over the network.

Client devices (e.g., data processing apparatus such as personal computers, smart phones, tablet computers, or laptop computers) can request map data for a geographic region of interest. The map data provided can be defined by a viewport, for example, which can be an element of the interactive mapping system graphical user interface (GUI). The viewport can be different shapes, e.g., rectangular or square, and can present map data of a particular geographic region. In response to the request for map data, one or more service providers can send the client device map data, which may be in the form of an image.

The client device then displays the map data or image in the viewport of the GUI (e.g., using a client web browser application). Map data can include map images (e.g., political or topographic map images), satellite images, business locations, popular landmarks, driving or walking directions, and vector graphics defining paths and regions overlaid on map images. Map data can also include various layers of related data, for example, a layer illustrating volcanoes in the Pacific Ocean or current traffic conditions.

FIG. 1 is a diagram of a graphical user interface of an example interactive mapping system 100. The interactive mapping system 100 contains a map image 110 showing a map of a portion of the earth's surface. The region displayed by the interactive mapping system 100 is defined by a viewport 140.

5 The interactive mapping system can include interface elements to control operation of the map, such as a panning control 120, a zoom control 130, a tilt control (not shown), or a rotation control (not shown). In some implementations, the user specifies a pan command by using an input device, e.g., a mouse, to drag the map image or manipulate the panning control 120. In some other implementations, the user specifies  
10 a pan command by dragging a finger across the screen of a touchscreen device. The interactive mapping system can provide data at multiple zoom levels (e.g., in response to a user input to the zoom control 130). Each subsequent zoom level provides more detail corresponding to a smaller geographic region.

Map data servers can provide images of map data in the form of tiles. Tiles are  
15 images that can be combined to form a larger, composite image. For example, a tile can be a 256x256 pixel image. Four such tiles can be combined to form a 512x512 pixel image. The map image 110, for example, can be broken up and provided as four separate tiles 141, 142, 143, and 144. Tile boundaries may or may not be visible in the viewport of the client device.

20 The tiles provided depend upon the current zoom level of the interactive mapping system. When zooming in using the zoom control 130, map data corresponding to a region of a previously provided tile can be subsequently provided as a composite of smaller, potentially higher resolution tiles. For example, map data for the region corresponding to tile 142 can be provided in a subsequent zoom level as tiles 151, 152,  
25 153, and 154.

Client devices can request tiles based on coordinates. The set of coordinates can be specified by the range of the user viewport. Coordinates can be latitude/longitude pairs or can be coordinates assigned by the interactive mapping system. Each tile can be  
30 referenced by a unique [x, y] pair of coordinates. For example, a client device can request tile 141 by the assigned coordinates [3, 4]. Tile 142 can be requested by assigned coordinates [4, 4]. Tile 143 can be requested by coordinates [3, 5], and tile 144 can be requested by coordinates [4, 5].

Particular tiles can be requested by including their coordinates in a URL of an HTTP request. For example, tile 141 can be requested by appending "?x=3&y=4" to a

resource locator (e.g. a URL) for the interactive mapping service. In mapping systems with multiple zoom levels, tiles can also be specified by a zoom level,  $z$ . Thus, at zoom level 3, a client device can request tile 141 with  $x=3$ ,  $y=4$ , and  $z=3$ , formulated as an HTTP request as “ $?x=3&y=4&z=3$ ”. Tile 151 (which is a smaller tile provided at a higher zoom level than tile 141) can be requested by specifying a higher zoom level, such as “ $?x=3&y=4&z=4$ ”.

As an interactive mapping system gathers more map data, tiles in the system can be updated. The updated information can, for example, reflect additional road information, higher resolution satellite imagery, or corrections to errors in existing map data. Tiles in the system can be assigned a version number to distinguish tiles with newer or older information. Client devices can request tiles of a specific version number,  $v$ , in an HTTP request. For example, version 2 of tile 141 at zoom level 3 could be requested by appending “ $?x=3&y=4&z=3&v=2$ ” to a resource locator (e.g. a URL) for the interactive mapping service.

FIG. 2 is a diagram of an example network environment 200 for serving map data. Serving map data over a network often involves communication between multiple servers in a series of requests. Map data ultimately served to a client device can be routed through multiple intermediary proxy servers or Internet service providers (ISPs). A proxy server is a server that mediates requests from clients to other servers in a network. An ISP provides client devices access to other servers on the network, which can be provided by a dial-up connection through a public switched telephone network, a digital subscriber line (DSL), cable broadband, WiFi, or any other network connection technology.

Specifically, the example network environment 200 includes map data servers 210. Map data servers 210 receive requests from client devices 242, 244, 246, and 248, respectively. The map data servers 210 serve map data back to the corresponding client devices. The requests and provided map data can be routed through proxy servers 222 and 224 and ISPs 232, 234, and 236 before reaching map data servers 210.

The provided map data can be cached by intermediary devices between the map data servers 210 and the client devices 242, 244, and 246. Client devices can also cache map data on a local storage device. Caching a resource on a network (e.g. map data) means that a device stores a local copy of the resource corresponding to a given resource locator and retrieves the copy of the resource on subsequent requests for the same network location instead of requesting the resource directly. Consequently, upon the next

request for the same resource locator (from the same or a different client device), the copy of the resource is served instead of requesting the resource again from the original server.

Caching network resources can reduce latency experienced by client devices by reducing the number of intermediate requests for a resource, and can also reduce load on upstream servers by reducing the number of requests for the original resource. In some implementations, a server will not subsequently modify a resource that has been identified as a cacheable resource. Cached resources that are subsequently modified introduce the possibility of client devices receiving inconsistent data.

In some implementations, a server identifies a resource that should be cached by including an appropriate header in an HTTP response providing the resource. The header includes a field indicating that servers forwarding the resource should store a copy of the resource and serve the copy on subsequent requests for the same resource. The header can also identify a time period after which the cached resource will expire, at which point the original resource should be requested again by the intermediate servers.

Alternatively, a server can indicate with the HTTP header that a resource should never be cached.

For example, map data servers 210 can indicate that a particular tile should be cached. When the tile is requested by a client device (e.g., client device 244), the tile will be provided to proxy server A 222, which will provide the tile to ISP B 234. The ISP B 234 will then provide the tile to the requesting client device. When proxy server A 222 receives the tile from map data servers 210, the proxy server A 222 reads the HTTP header and determines that the tile should be cached. Proxy server A 222 then creates a local copy of the tile to be served on subsequent requests for that tile. For example, if proxy server A 222 receives a subsequent request for the same tile, proxy server A 222 responds by serving the stored local copy of the tile rather than requesting the tile from map data servers 210. ISPs 232, 234, and 236 cache the tile in the same way by reading the HTTP header of received tiles. Client devices 242, 244, 246, and 248 can also cache a local copy of the tile, which will be read from a local storage device rather than requesting the tile from their respective ISPs.

“Cache hits” refer to instances of a client device or an intermediate server finding a locally stored copy of a cached resource. “Cache misses,” on the other hand, refer to instances where no locally stored copy of a resource is found on the client device or on any of the intermediate servers, or when a locally stored copy has expired or is

otherwise invalid. Cache misses require requesting the original resource from the original server, e.g., map data servers 210.

Because cache hits reduce latency and decrease load on servers, an interactive mapping system can attempt to maximize the number of cache hits and minimize the number of cache misses on requests for resources such as tiles. To improve caching performance, an interactive mapping system can implement a separate directory structure used for serving and caching tiles.

FIG. 3 is a diagram of an example directory structure 300 for caching tiles. The example directory structure is implemented as a quadtree, in which each node of the directory structure has four child nodes. Tiles in the interactive mapping system can also be organized in a quadtree structure, but the example directory structure for caching shown in FIG. 3 is not necessarily related to the structure of the interactive mapping system and can be implemented as an entirely separate structure. Furthermore, the directory structure in FIG. 3 could be implemented as another kind of tree, e.g., as a B-tree.

The leaf nodes shown in FIG. 3 (e.g., nodes 311-314, 321-324, 331-334, and 341-344) correspond to individual tile versions in the interactive mapping system. Each leaf node contains information required to retrieve the corresponding tile. For example, a leaf node could contain the information,  $[x=2, y=3, z=2, v=2]$ , for map coordinates  $x$  and  $y$ , zoom level  $z$ , and tile version  $v$ , which could be used to request a tile with a URL containing “ $?x=2&y=3&z=2&v=2$ ”. In response to a URL request for this tile, the interactive mapping system can provide the corresponding tile data. In addition, leaf nodes of the directory structure can contain per-tile data (e.g., locations of businesses within the tile region). For brevity, however, only the version number of each tile is shown in FIG. 3.

In some implementations, a separate directory structure is generated for each zoom level of the interactive mapping system. Thus, all tiles corresponding to the leaf nodes of the example directory structure shown in FIG. 3 are at the same zoom level in the interactive mapping system. The example directory structure contains only 16 leaf nodes, and therefore only two levels. However, an interactive mapping system can contain millions of tiles at a given zoom level, and thus the directory structure would accordingly contain more levels than the example directory structure shown in FIG. 3.

The intermediate nodes 310, 320, 330, and 340 contain a hash of the contents of their respective child nodes. For example, node 310 can contain a hash of the

concatenation of the contents of leaf nodes 311, 312, 313, and 314. A hash is a string of characters generated by a hash function. A hash function converts input data into a hash, which is a sequence of hash characters. Each hash character can correspond to a bit string and can be represented in various character encodings, such as hexadecimal or Base64.

5 Similarly, the root node 350 contains a hash of the concatenation of its child nodes, nodes 310, 320, 330, and 340.

In some implementations, hashes are also used to assign identifying URLs for map tiles. When the identifying URLs are maintained in a directory structure such as the one shown in FIG. 3, the URLs of the map tiles do not need to be predictable (e.g. “x=3&y=4” for [x=3, y=4]) and can instead be arbitrary. The URL of each map tile can then be generated by a hash of the image data in each map tile, instead of a predictable concatenation of coordinates, version number, and zoom level. When URLs are thus generated by a hash of image data, visually identical tiles (e.g., solid color tiles for oceans, uninhabited regions, or regions for which data is unavailable) are assigned identical URLs, which further increases cache hits. Visually identical tiles particularly increase cache hits on the client device itself, eliminating HTTP requests to a server.

The contents of each node in the directory structure can be used to access that node as a network resource location. For example, the hash contained in node 310 can be used as a URL for a client device to access node 310. The URL request for node 310 can be a URL that includes “?id=26fb6”, where “id” represents the hash identifier. In response to a request for an intermediate node, the interactive mapping system can provide a list of that node’s child nodes. For example, in response to a URL request for node 310, the interactive mapping system could provide identifying information for the child nodes of node 310, which are leaf nodes 311, 312, 313, and 314. The identifying information can include the x and y coordinates, the zoom level, and the version number of each respective tile.

When the contents of a child node changes, the contents of the child node’s parent node will also change (because each intermediate node contains a hash of a concatenation of the contents of its child nodes). Therefore, changes in the directory structure get propagated up all the way to the root node. In other words, the root node changes whenever any child node in the directory structure changes.

In some implementations, because the leaf nodes of the directory structure are cacheable resources (e.g. map tiles), the contents of individual child nodes do not change. Instead, a new child node is created and associated with an appropriate parent node.

Therefore, requests for an old map tile can still result in access to the old map tile, even after the map tile has been updated. Accessing the root node thus provides a snapshot of tiles of the world because only branches of the directory structure reachable from the accessed root node will be subsequently traversed by a client. Newly added nodes are  
5 reachable only after re-requesting the root node.

When a client device loads the user interface of the interactive mapping system (e.g. the user interface as shown in FIG. 1), the viewport identifies the map tiles that should be loaded. In this situation, the interactive mapping system can immediately provide identifying information for URLs of the most recently updated tiles, such as the x  
10 and y coordinates, zoom level, and version number, as well as URLs for all ancestor nodes. The tiles are served through intermediate proxy servers and ISPs and are cached, and subsequent requests for the cached tiles result in cache hits.

However, in some implementations (e.g. when a client device does not indicate a region of interest by a viewport), the interactive mapping system must be interrogated in  
15 order for the client device to obtain identifying information about which version of map tiles should be requested.

FIG. 4 is a sequence diagram of an example client interaction with the interactive mapping system. The client device interacts with the interactive mapping system through a proxy server to obtain identifying information about which version of map tiles should  
20 be requested.

The client device 410 requests the root of the directory structure for caching tiles 402. In some scenarios where a separate directory structure is maintained for each zoom level, the client device specifies a zoom level in its request for the root. The proxy server 420 forwards the request 404 to the map data server 430. In response to the request, the  
25 map data server 430 provides a list of the root's child nodes 406, which the proxy server 420 forwards 408 to the client device 410.

For example, in response to a request for the root node, the client device receives a list of child nodes  $[[26fb6, b7f03], [c7090, f1038]]$  when requesting a tile at  $x=1$  and  $y=2$ . The client can specify which of the four child nodes should be subsequently  
30 requested by a pair of indices, e.g.,  $[0, 1]$  or  $[1, 1]$ . To identify which child of the root node should be subsequently requested, the client device uses the x and y coordinates of the currently requested tile. In some implementations, bits in the x and y coordinates identify the appropriate child of the root node. The identifying bits correspond to the level of the current node in the directory structure. For example, because the root node

is at the first level of the directory structure, the first bits (i.e. the most significant bits) of coordinates  $x=1$  (binary “01”) and  $y=2$  (binary “10”) identify the child that should be subsequently requested, yielding 0 and 1 respectively. Therefore, the client device would next request the node identified by [0, 1] in the list of returned child nodes, which is  
5 “b7f03”.

The client device 410 makes a request for a node 412 using the node identifier “b7f03”. In some implementations, the request is a URL generated by appending the node identifier to the map service URL (e.g. “http://example.com/map?node=b7f03”). The proxy server 420 forwards the request 414 to the map data server 430. In response to the  
10 request, the map data server 430 provides a list of the node’s children 416, which the proxy server 420 forwards 418 to the client device 410.

The client device again uses the  $x$  and  $y$  coordinates of the tile being requested to identify which of the child nodes should be requested. In this example, the children of the requested nodes are leaf nodes containing identifying information for map tiles. For  
15 example, the client device could receive a list of nodes  $[[x=0&y=2&z=2&v=2, x=0&y=3&z=2&v=2], [x=1&y=2&z=2&v=2, x=1&y=3&z=2&v=2]]$ . The client device identifies the appropriate tile identifier information using the  $x$  and  $y$  coordinates being requested. Because these are leaf nodes, the last (e.g. least significant) bits of the tile coordinates are used to identify the appropriate tile. In this case, the least significant bits  
20 of  $x=1$  (binary “01”) and  $y=2$  (binary “10”) are 1 and 0 respectively, so the node with identifier “ $x=1&y=2&z=2&v=2$ ” is identified.

The client device makes a request for the map tile 422 using the tile identifier “ $x=1&y=2&z=2&v=2$ ”. In some implementations, the request is a URL generated by appending the tile identifier to the map service URL (e.g. “http://example.com/map?  
25  $x=1&y=2&z=2&v=2$ ”). The proxy server 420 forwards the request 424 to the map data server 430. In response to the request, the map data server 430 provides the tile data 426, which the proxy server 420 forwards 428 to the client device 410.

The map data server 430 can include appropriate headers with responses so that map resources requested by clients are cached. After the proxy server 420 receives tile  
30 data 426, subsequent requests for the same URL will result in cache hits. If the client device 410 requests the same tile URL 432, the proxy server will respond with a cached copy of the tile data 434. In some implementations, the tile data is also cached on the client device 410, and thus the client device can retrieve a cached copy of the tile data without requesting the tile data from the proxy server 420.

Similarly, if the client device 410 requests the URL of a previously requested node 442, the proxy server will respond with a cached copy of the node data 444. In some implementations, only the root node is never cached. Cached resources can be served to multiple different clients, so if another client different from client 410 requested the URL of a previously requested node (e.g. request 442), the proxy server 420 would respond with a cached rather than requesting the original data from the map data server 430.

FIG. 5 is a diagram of an example update of the directory structure for caching tiles. Because each intermediate node contains a hash of the contents of its respective child nodes, updates to child nodes are propagated up the directory structure, changing each ancestor along the way. Updates to any child node will thus update the root node.

The tile corresponding to leaf node 514 is updated in the interactive mapping system. A new leaf node 515 is created, corresponding to version 3 of the tile. As a result of the update, a new parent node 516 is created, containing a hash of the contents of its child nodes 511, 512, 513, and new node 515. As a result of new parent node 516, a new root node 555 is created, containing a hash of the contents of its child nodes 516, 520, 530, and 540.

In some implementations, the old root 550 and old parent node 510 are still accessible by their URLs for a specified time period after the new root 555 is created. However, the map data server immediately identifies new root 555 in response to requests for the root node. After a specified time period has passed, the map data server can carry out a garbage collection process that will erase root node 550 and node 510. However, if node 510 has been cached on an intermediate proxy server, requests by client devices for this node will continue to generate cache hits until the root node is requested again.

FIG. 6 is a sequence diagram of an example client interaction with the interactive mapping system after a tile is updated. Until the client device re-requests the root node, the client device will continue to receive cached copies of requested tiles.

New tile data becomes available and a new version of a tile is created 602, resulting in a new parent node being created 604, and a new root node being created 606. In some implementations, new tiles versions are added rather than new tiles replacing old tiles.

The client device requests a tile at a specified set of coordinates 612. Though a new version of the tile at the specified coordinates is available, the client device receives a cached copy of old tile data 614. Requesting client devices continue to receive the

cached version of the old tile until a request for the root node is received from a client device.

In some implementations, the client device re-requests the root node when a new session is started 616. The root node can also be re-requested if the cached entries on proxy server 620 expire. After requesting the root node, the client device will send a series of requests to traverse the directory structure in order to obtain a resource locator for a requested tile.

The client device requests the root node 622 in connection with a requested tile located at coordinates  $[x, y]$ . The proxy server 620 forwards the request 624 to the map data server 630, and the map data server 630 responds with a list of the new root's child nodes 636, which the proxy server 620 forwards 638 to the client device 610. The client identifies which of the root's child nodes should be subsequently requested by using the most significant bits of the  $x$  and  $y$  coordinates. Among the new root's child nodes will be a resource locator for a new intermediate node created after and in response to the addition of a node for the new tile.

The client requests the new node 632. The proxy server 620 forwards the request to the map data server 630. In response to the request, the map data server 630, provides a list of the node's children 636, which the proxy server 620 forwards 648 to the client device 610. The client device again uses the  $x$  and  $y$  coordinates of the tile being requested to identify which of the child nodes should be requested. The client device makes a request for the tile 642, which the proxy server 620 forwards 644 to the map data server 630. In response to the request, map data server 630 provides the new tile data 646, which the proxy server 620 forwards 648 to the client device 610. Subsequent requests for the tile at coordinates  $[x, y]$  (e.g. request 652) will result in the proxy server 620 responding with a cached copy of new tile data 654.

Embodiments of the subject matter and the operations described in this specification can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions, encoded on computer storage medium for execution by, or to control the operation of, data processing apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated

electrical, optical, or electromagnetic signal that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. A computer storage medium can be, or be included in, a computer-readable storage device, a computer-readable storage substrate, a random or serial access memory array or  
5 device, or a combination of one or more of them. Moreover, while a computer storage medium is not a propagated signal, a computer storage medium can be a source or destination of computer program instructions encoded in an artificially-generated propagated signal. The computer storage medium can also be, or be included in, one or more separate physical components or media (e.g., multiple CDs, disks, or other storage  
10 devices).

The operations described in this specification can be implemented as operations performed by a data processing apparatus on data stored on one or more computer-readable storage devices or received from other sources.

The term “data processing apparatus” encompasses all kinds of apparatus, devices,  
15 and machines for processing data, including by way of example a programmable processor, a computer, a system on a chip, or multiple ones, or combinations, of the foregoing. The apparatus can include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can also include, in addition to hardware, code that creates an execution  
20 environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, a cross-platform runtime environment, a virtual machine, or a combination of one or more of them. The apparatus and execution environment can realize various different computing model infrastructures, such as web services, distributed computing and grid computing  
25 infrastructures.

A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine,  
30 object, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language document), in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub-programs, or

portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform actions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital 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 performing actions in accordance with instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device (e.g., a universal serial bus (USB) flash drive), to name just a few. Devices suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user

can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for  
5 example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that  
10 includes a front-end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a  
15 communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), an inter-network (e.g., the Internet), and peer-to-peer networks (e.g., ad hoc peer-to-peer networks).

The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication  
20 network. 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. In some embodiments, a server transmits data (e.g., an HTML page) to a client device (e.g., for purposes of displaying data to and receiving user input from a user interacting with the client device). Data generated at the client device (e.g., a result of the user  
25 interaction) can be received from the client device at the server.

While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any inventions or of what may be claimed, but rather as descriptions of features specific to particular embodiments of particular inventions. Certain features that are described in this specification in the  
30 context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a

claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

5 Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to  
10 achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and  
15 systems can generally be integrated together in a single software product or packaged into multiple software products.

Thus, particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable  
15 results. In addition, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

What is claimed is:

CLAIMS

- 5 1. A method for caching tiles in an interactive mapping system comprising:  
receiving, from a first client, a request for a root node of a directory structure in  
which resource locators for nodes are generated based on a hash of resource locators of  
respective one or more descendant nodes, wherein requests for parent nodes generate  
responses containing resource locators of respective one or more descendant nodes, and  
10 wherein leaf nodes are associated with corresponding resource locators for tiles in an  
interactive mapping system;  
serving, to the first client, a first configuration of tiles, wherein each intermediate  
node and each tile is served as a cacheable resource;  
receiving an indication of a new node added to the directory structure, the new  
15 node corresponding to a new version of a tile;  
adding to the directory structure one or more new ancestor nodes of the new node;  
receiving, from a second client, a request for the root node; and  
serving, to the second client, a second configuration of tiles including the new  
node, while continuing to serve the first configuration of tiles to one or more other clients  
20 that requested the root node before the one or more new ancestor nodes of the new node  
were added.
2. The method of claim 1, wherein adding one or more new ancestor nodes of the  
new node comprises switching from serving the first configuration of tiles to serving the  
25 second configuration of tiles after a new root node is added.
3. The method of claim 2, wherein switching to serving the second configuration of  
tiles comprises swapping a root node indicator from the root node to the new root node.
- 30 4. The method of claim 3, further comprising serving, in response to receiving, from  
a second client, a request for the root node, a response containing resource locators of one  
or more descendant nodes of the new root node.
5. The method of claim 1, where the resource locators for tiles include a version.

35

6. The method of claim 1, further comprising serving the second configuration of tiles with an indication that one or more tiles are cacheable.
7. The method of claim 1, where updating resource locators for each of one or more ancestor nodes of the new node comprises adding to the directory structure the one or more ancestor nodes for the new node.
8. The method of claim 1, where the resource locators for tiles include map coordinates.
9. The method of claim 1, where the resource locators for tiles include a hash of tile data.
10. The method of claim 1, where resource locators for the root node and each parent node are generated by a hash function.
11. The method of claim 10, where resource locators for the root node and each parent node are generated by a hash of a concatenation of resource locators of one or more respective descendant nodes.
12. The method of claim 1, further comprising generating a separate directory structure for each of one or more zoom levels in the interactive mapping system.
13. The method of claim 1, where the directory structure is a quadtree or a B-tree.
14. The method of claim 1, where the indication that the tile is cacheable comprises an HTTP header.
15. A system comprising:  
one or more computers; and  
a computer-readable storage device storing instructions that, when executed by the one or more computers, cause the one or more computers to perform operations comprising:  
receiving, from a first client, a request for a root node of a directory

structure in which resource locators for nodes are generated based on a hash of resource locators of respective one or more descendant nodes, wherein requests for parent nodes generate responses containing resource locators of respective one or more descendant nodes, and wherein leaf nodes are associated with corresponding resource locators for tiles in an interactive mapping system;

5                   serving, to the first client, a first configuration of tiles, wherein each intermediate node and each tile is served as a cacheable resource;

                  receiving an indication of a new node added to the directory structure, the new node corresponding to a new version of a tile;

10                   adding to the directory structure one or more new ancestor nodes of the new node;

                  receiving, from a second client, a request for the root node; and

                  serving, to the second client, a second configuration of tiles including the new node, while continuing to serve the first configuration of tiles to one or more other

15                   clients that requested the root node before the one or more new ancestor nodes of the new node were added.

16.     The system of claim 15, wherein adding one or more new ancestor nodes of the new node comprises switching from serving the first configuration of tiles to serving the

20                   second configuration of tiles after a new root node is added.

17.     The system of claim 16, wherein switching to serving the second configuration of tiles comprises swapping a root node indicator from the root node to the new root node.

25     18.     The system of claim 17, where the operations further comprise serving, in response to receiving, from a second client, a request for the root node, a response containing resource locators of one or more descendant nodes of the new root node.

19.     The system of claim 15, where the resource locators for tiles include a version.

30

20.     The system of claim 15, where the operations further comprise serving the second configuration of tiles with an indication that one or more tiles are cacheable.

21. The system of claim 15, where updating resource locators for each of one or more ancestor nodes of the new node comprises adding to the directory structure the one or more ancestor nodes for the new node.
- 5 22. The system of claim 15, where the resource locators for tiles include map coordinates.
23. The system of claim 15, where the resource locators for tiles include a hash of tile data.
- 10 24. The system of claim 15, where resource locators for the root node and each parent node are generated by a hash function.
25. The system of claim 24, where resource locators for the root node and each parent node are generated by a hash of a concatenation of resource locators of one or more respective descendant nodes.
- 15 26. The system of claim 15, where the operations further comprise generating a separate directory structure for each of one or more zoom levels in the interactive mapping system.
- 20 27. The system of claim 15, where the directory structure is a quadtree or a B-tree.
28. The system of claim 15, where the indication that the tile is cacheable comprises an HTTP header.
- 25

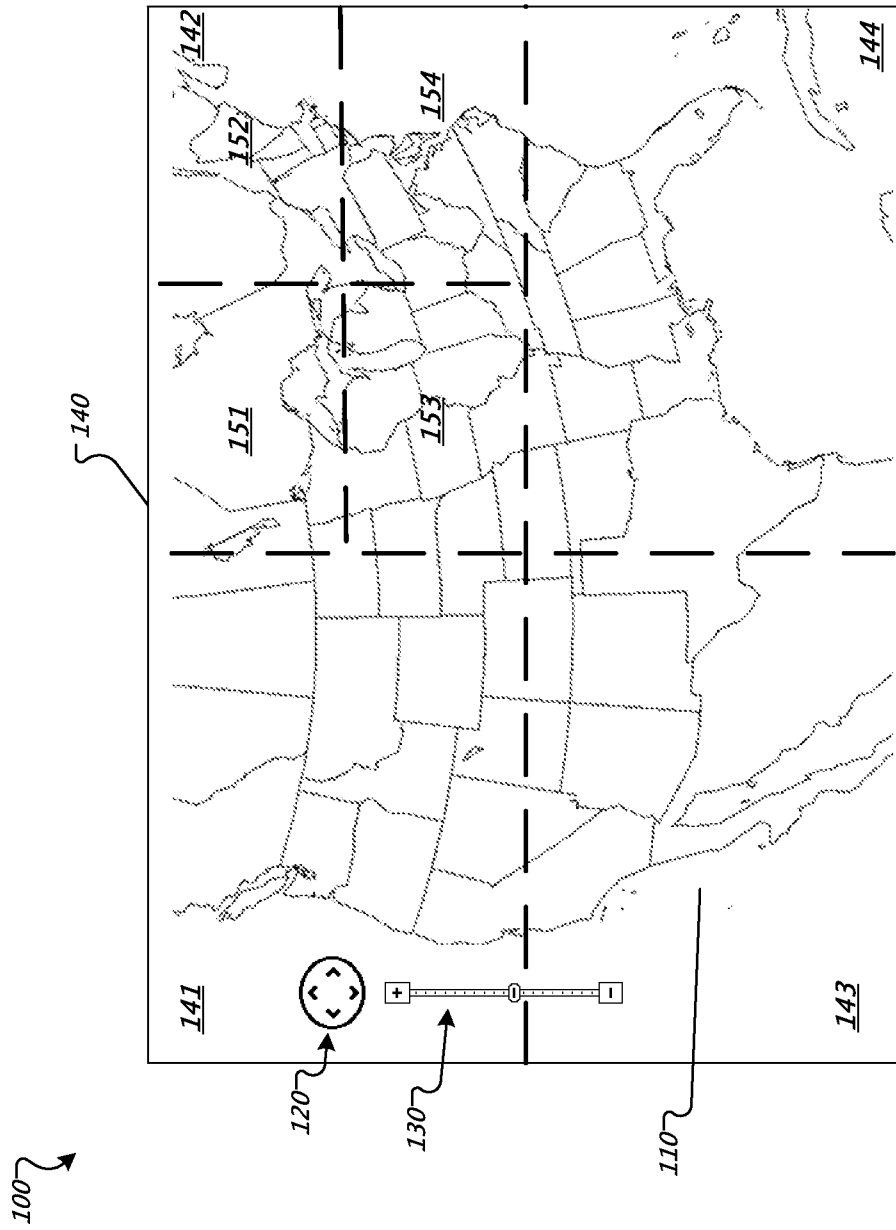


FIG. 1




FIG. 2

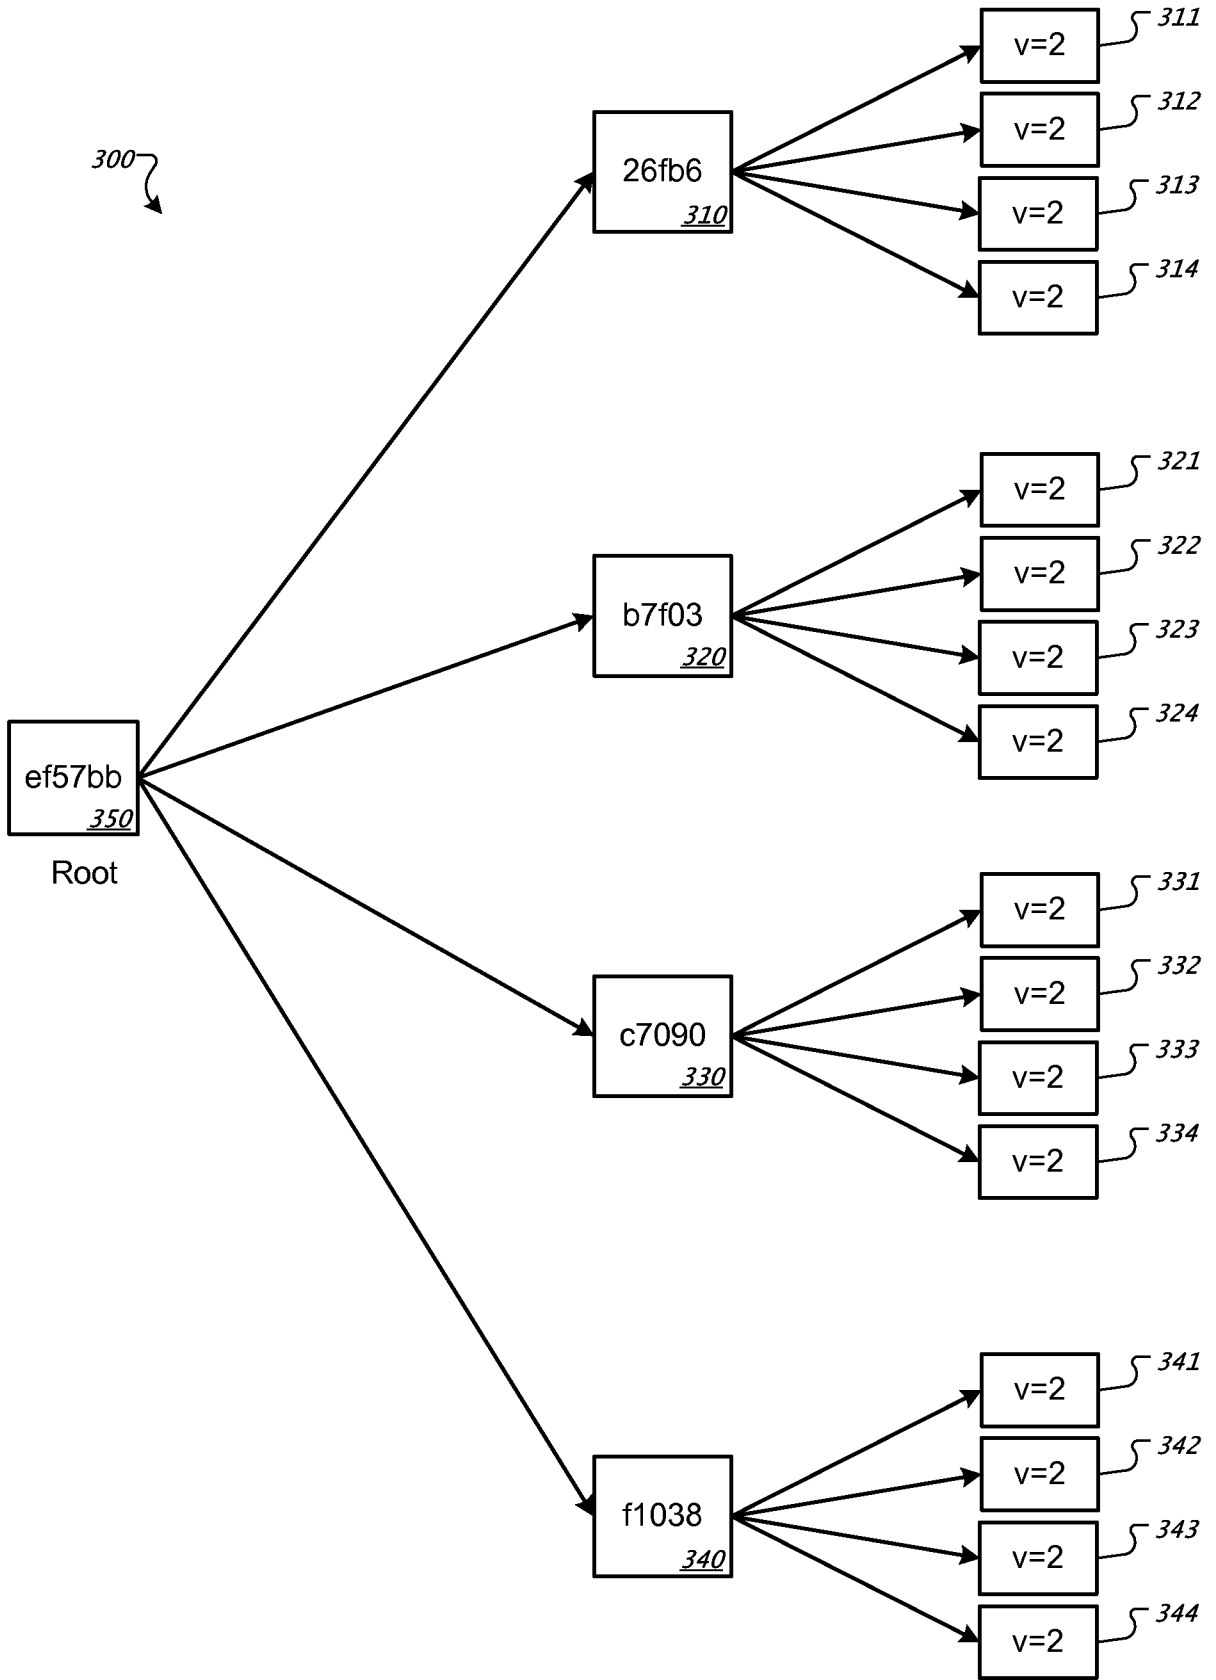


FIG. 3

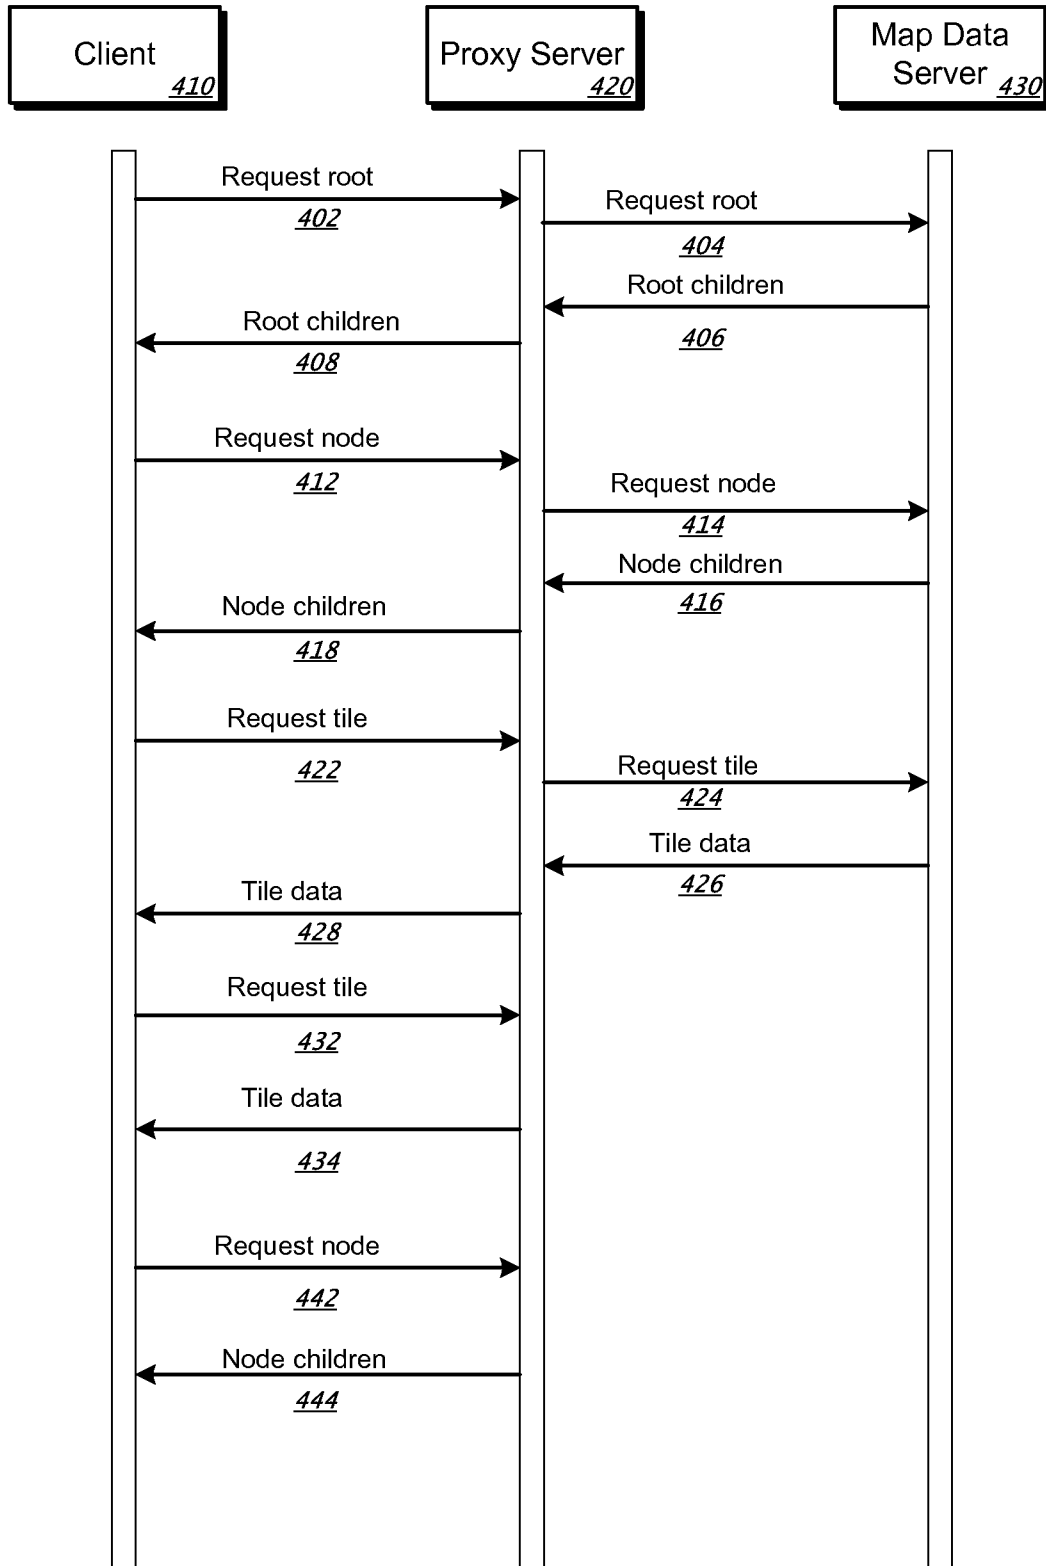


FIG. 4

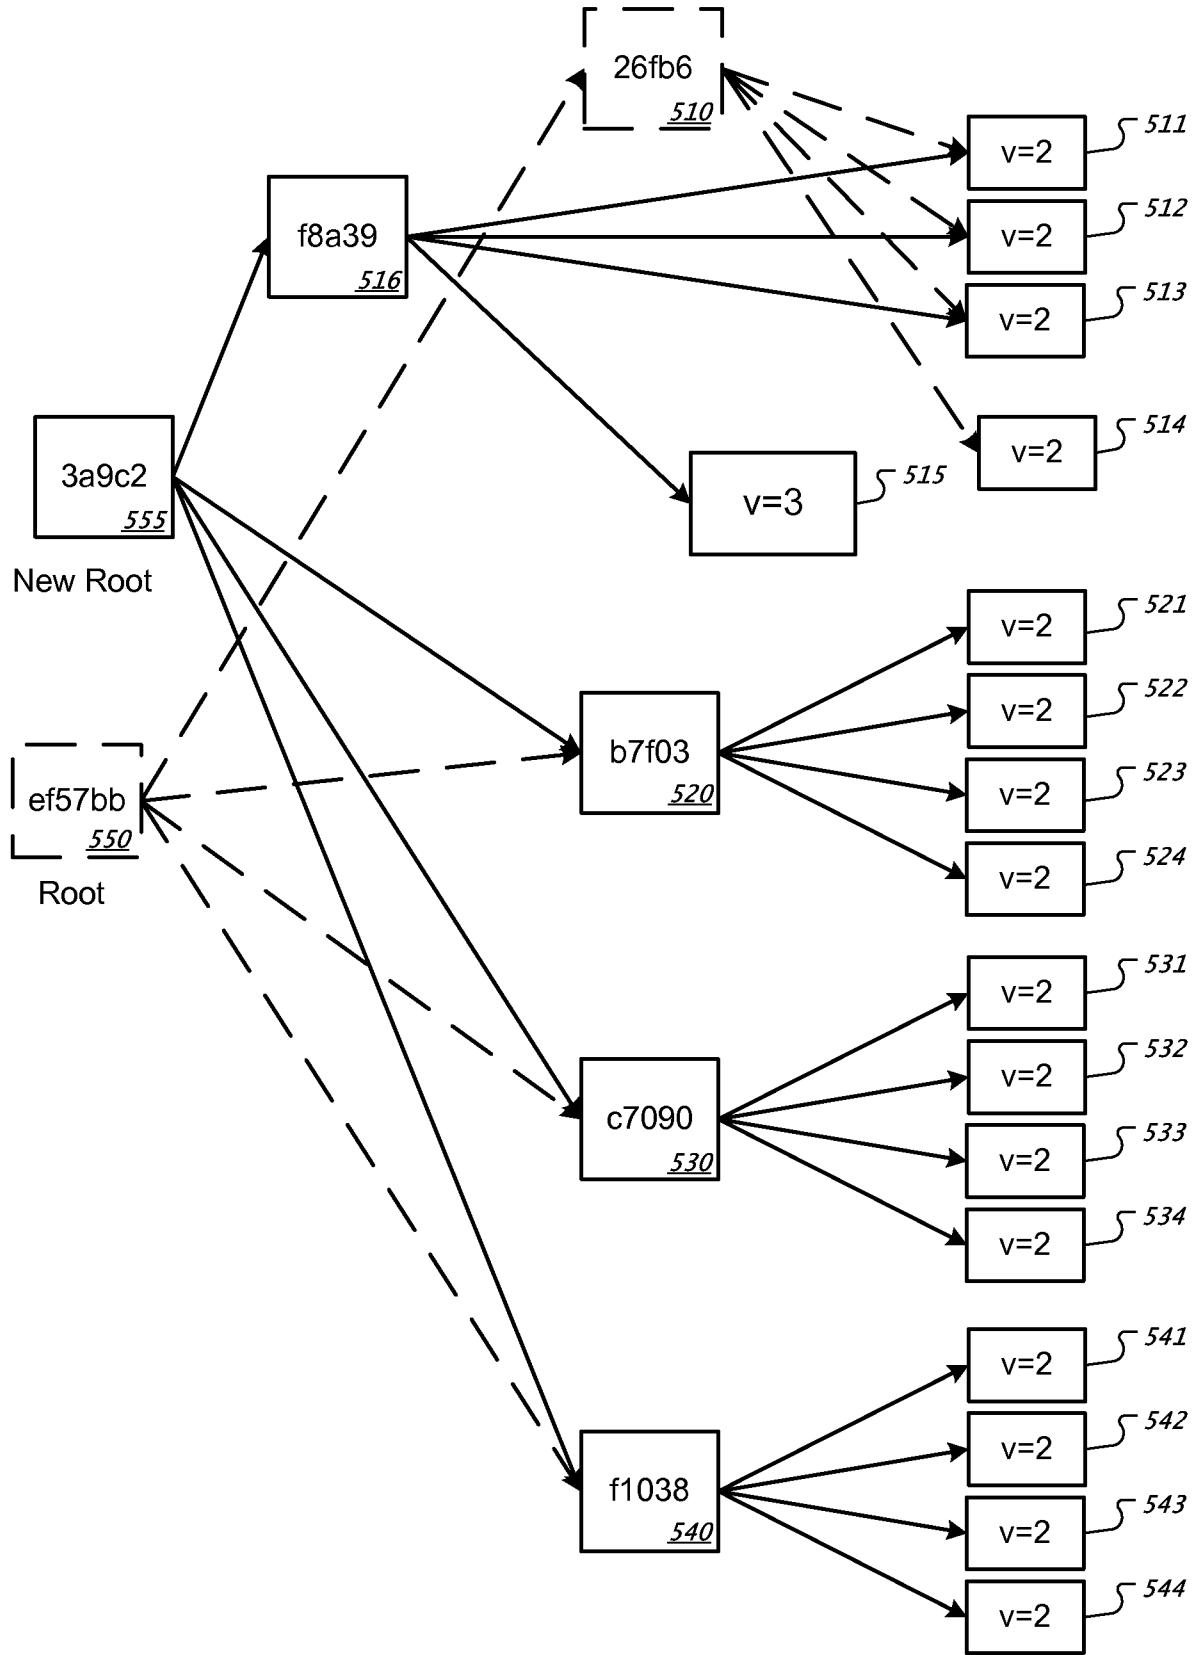


FIG. 5

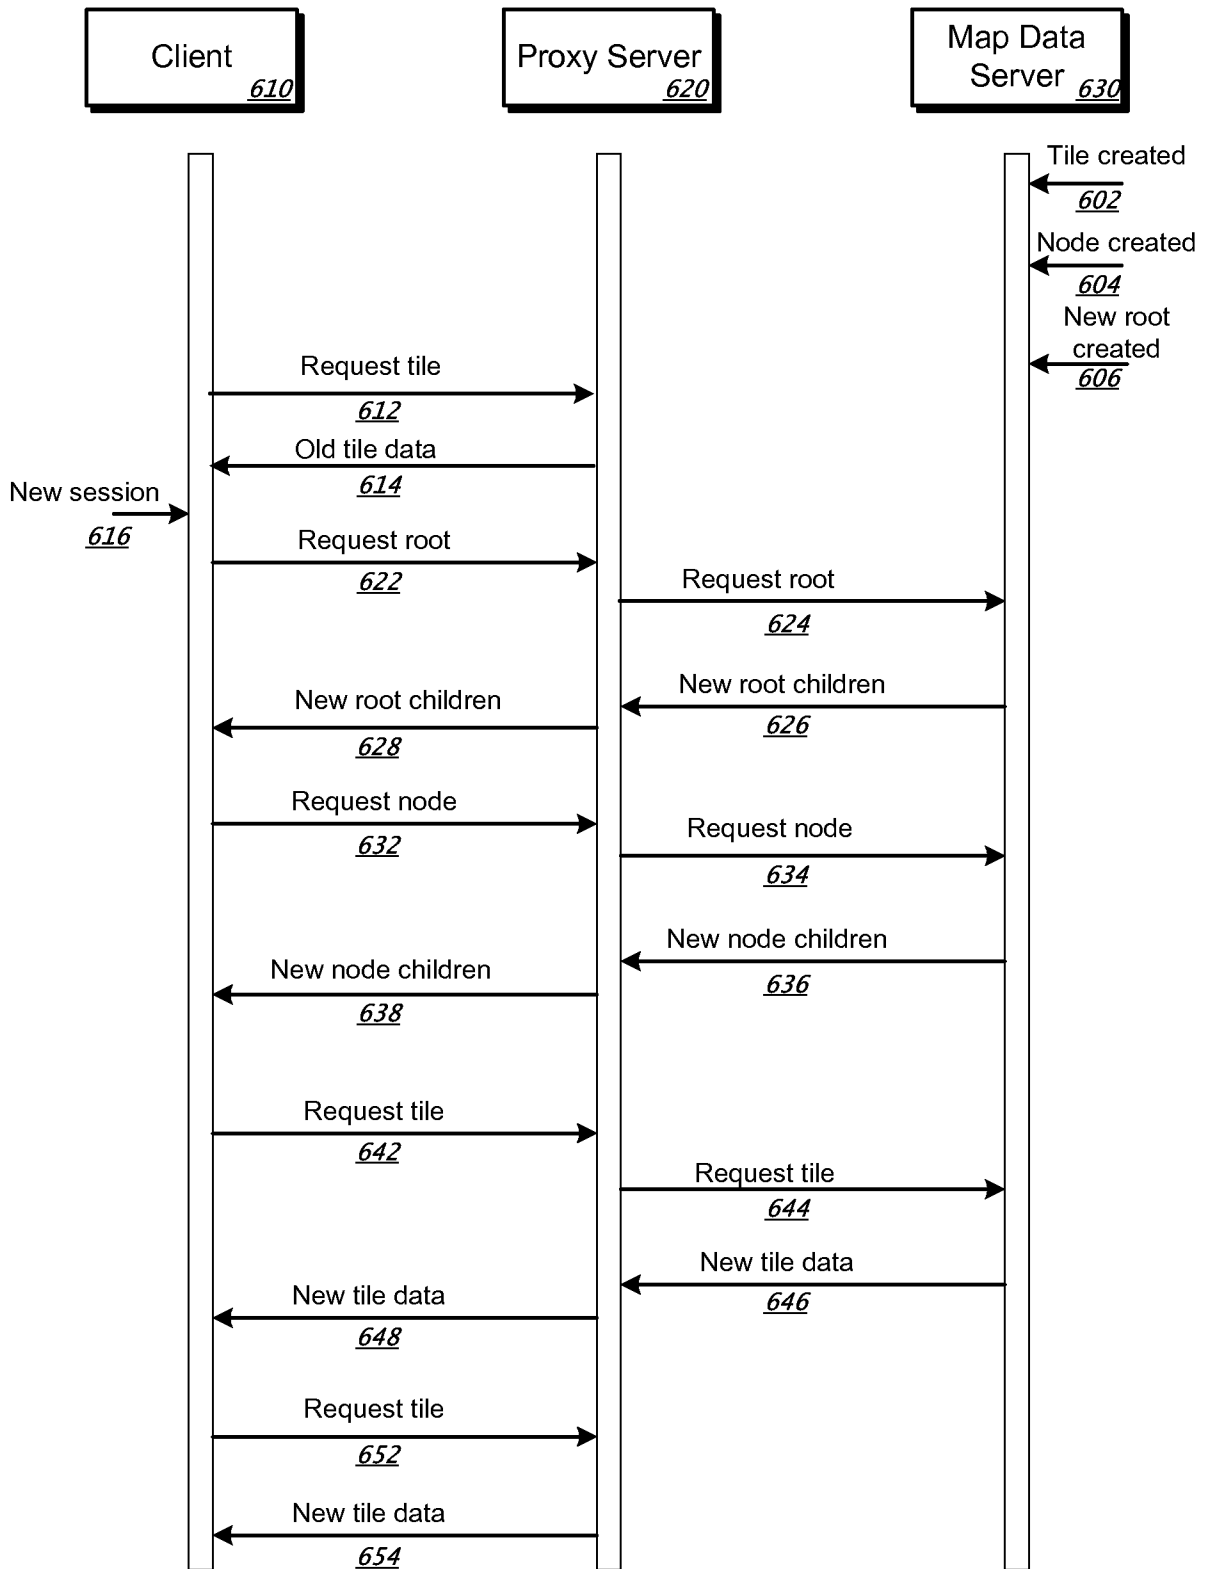




FIG. 6

## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2012/022577**

<b>A. CLASSIFICATION OF SUBJECT MATTER</b>		
<i>G06F 17/40(2006.01)i, G06F 15/16(2006.01)i</i>		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols) G06F 17/40; G06F 15/173; H04L 29/06; G06F 9/46; G06F 12/00		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched Korean utility models and applications for utility models Japanese utility models and applications for utility models		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) eKOMPASS(KIPO internal) & Keywords: caching, resource, map, directory and request.		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 97-06485 A1 (APPLE COMPUTER, INC.) 20 February 1997 See abstract; pages 1-8; claims 1-13; and figures 1-4.	1-28
A	EP 0618708 A2 (INTERNATIONAL BUSINESS MACHINES) 05 October 1994 See abstract; columns 1-12; claims 1-11; and figures 1-5.	1-28
A	US 2008-0209040 A1 (RATHI PRADEEP KUMAR) 28 August 2008 See abstract; pages 1-5; claims 1-25; and figures 1-4.	1-28
A	US 2007-0079072 A1 (COLLIER JOSH D. et al.) 05 April 2007 See abstract; pages 1-9; claims 1-20; and figures 1-6.	1-28
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family		
Date of the actual completion of the international search 30 APRIL 2012 (30.04.2012)		Date of mailing of the international search report <b>04 MAY 2012 (04.05.2012)</b>
Name and mailing address of the ISA/KR  Korean Intellectual Property Office Government Complex-Daejeon, 189 Cheongsu-ro, Seo-gu, Daejeon 302-701, Republic of Korea Facsimile No. 82-42-472-7140		Authorized officer Ahn Ji Hyun Telephone No. 82-42-481-5767 

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2012/022577**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 97-06485 A1	20.02.1997	AU 6951896 A EP 0843850 A1 US 05778383A A	05.03.1997 27.05.1998 07.07.1998
EP 0618708 A2	05.10.1994	JP 02-511644 B2 JP 06-309264 A JP 2511644 B2 US 05511208A A	16.04.1996 04.11.1994 03.07.1996 23.04.1996
US 2008-0209040 A1	28.08.2008	US 7844710 B2	30.11.2010
US 2007-0079072 A1	05.04.2007	None	