



(51) International Patent Classification:  
*G06F 17/30* (2006.01)

(21) International Application Number:  
PCT/US2014/029724

(22) International Filing Date:  
14 March 2014 (14.03.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
61/799,817 15 March 2013 (15.03.2013) US  
61/799,986 15 March 2013 (15.03.2013) US  
61/800,036 15 March 2013 (15.03.2013) US  
61/799,846 15 March 2013 (15.03.2013) US  
61/799,131 15 March 2013 (15.03.2013) US

(72) Inventors; and

(71) Applicants : RANA, Ahad [US/US]; 2424 Banyan Drive, Los Angeles, CA 90049 (US). SHIMANOVSKY, Boris [US/US]; 2417 Bagley Avenue, Los Angeles, CA 90034 (US).

(74) Agents: GOLDENBERG, Richard, A. et al.; Wilmer Cutler Pickering Hale and Dorr LLP, 60 State Street, Boston, MA 02109 (US).


(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))



WO 2014/145069 A1

(54) Title: APPARATUS, SYSTEMS, AND METHODS FOR PROVIDING LOCATION INFORMATION

(57) Abstract: The disclosed apparatus, systems, and methods relate to a location query mechanism that can efficiently determine whether a target entity is located within a region of interest (ROI). At a high level, the location query mechanism can be configured to represent a ROI using one or more polygons. The location query mechanism can, in turn, divide (e.g., tessellate) the one or more polygons into sub-polygons. Subsequently, the location query mechanism can use the sub-polygons to build an index system that can efficiently determine whether a particular location is within any of the sub-polygons. Therefore, when a computing device queries whether a particular location is within the region of interest, the location query mechanism can use the index system to determine whether the particular location is within any of the sub-polygons.

**APPARATUS, SYSTEMS, AND METHODS**  
**FOR PROVIDING LOCATION INFORMATION**

Cross reference to related applications

[0001] This application claims benefit of the earlier filing date, under 35 U.S.C. §119(e), of:

- U.S. Provisional Application No. 61/799,986, filed on March 15, 2013, entitled “SYSTEM FOR ANALYZING AND USING LOCATION BASED BEHAVIOR”;
- U.S. Provisional Application No. 61/800,036, filed on March 15, 2013, entitled “GEOGRAPHIC LOCATION DESCRIPTOR AND LINKER”;
- U.S. Provisional Application No. 61/799,131, filed on March 15, 2013, entitled “SYSTEM AND METHOD FOR CROWD SOURCING DOMAIN SPECIFIC INTELLIGENCE”;
- U.S. Provisional Application No. 61/799,846, filed March 15, 2013, entitled “SYSTEM WITH BATCH AND REAL TIME DATA PROCESSING”; and
- U.S. Provisional Application No. 61/799,817, filed on March 15, 2013, entitled “SYSTEM FOR ASSIGNING SCORES TO LOCATION ENTITIES”.

[0002] This application is also related to:

- U.S. Patent Application No. 14/214,208, entitled “APPARATUS, SYSTEMS, AND METHODS FOR ANALYZING MOVEMENTS OF TARGET ENTITIES,” filed on the even-date herewith;
- U.S. Patent Application No. 14/214,296, entitled “APPARATUS, SYSTEMS, AND METHODS FOR PROVIDING LOCATION INFORMATION,” filed on the even-date herewith;
- U.S. Patent Application No. 14/214,213, entitled “APPARATUS, SYSTEMS, AND METHODS FOR CROWDSOURCING DOMAIN SPECIFIC INTELLIGENCE,” filed on the even-date herewith;
- U.S. Patent Application No. 14/214,219, entitled “APPARATUS, SYSTEMS, AND METHODS FOR BATCH AND REALTIME DATA PROCESSING,” filed on the even-date herewith;
- U.S. Patent Application No. 14/214,309, entitled “APPARATUS, SYSTEMS, AND METHODS FOR ANALYZING CHARACTERISTICS OF ENTITIES OF INTEREST,” filed on the even-date herewith; and

- U.S. Patent Application No. 14/214,231, entitled “APPARATUS, SYSTEMS, AND METHODS FOR GROUPING DATA RECORDS,” filed on the even-date herewith.

[0003] The entire content of each of the above-referenced applications (including both the provisional applications and the non-provisional applications) is herein incorporated by reference.

#### Field of the Invention

[0004] The present disclosure generally relates to data processing systems, and specifically, to data processing systems that can provide information on geographic location of an entity.

#### Background

[0005] As location-aware devices, such as Global Positioning System (GPS)-enabled mobile phones, have become popular over the years, it is increasingly desirable to quickly and efficiently determine whether a device is within a region of interest. Similarly, developers of sports-related mobile applications may wish to provide a different interface for users when users are inside a football stadium or tailgating in the stadium parking lot. Additionally, processing data with location information from, for example, user logs, location tagged social network information (e.g. a stream of tweets from twitter), or similar data can benefit from additional contextual information such as whether the stated location is in a region of interest such as a shopping mall. In such cases, it is desirable to a) define region(s) of interest, b) determine whether a geographic location point (e.g. latitude and longitude) is within the region(s) of interest c) if so, identify the region corresponding to the geographic location point and d) based on the identified region, if any, determine an action to be performed based on the identified region, such as a particular advertisement, user interface, or other computer logic that should occur based on the location being inside the region of interest.

#### Summary

[0006] In general, in an aspect, embodiments of the disclosed subject matter can include an apparatus. The apparatus includes a processor configured to run one or more modules stored in memory. The one or more modules are configured to receive one or more polygons associated with a region of interest, determine a plurality of sub-polygons that are contained within the one or more polygons, wherein each of the sub-polygons is associated with a unique code, and generate a first index system based on at least a subset of the plurality of

sub-polygons, thereby providing an efficient mechanism to determine whether a particular location is within the region of interest.

[0007] In general, in an aspect, embodiments of the disclosed subject matter can include a method. The method includes receiving, at an index generation module of a computing system, one or more polygons associated with the region of interest, determining, at the index generation module, a plurality of sub-polygons that are contained within the one or more polygons, wherein each of the sub-polygons is associated with a unique code, and generating, at the index generation module, a first index system based on at least a subset of the plurality of sub-polygons, thereby providing an efficient mechanism to determine whether a particular location is within the region of interest.

[0008] In general, in an aspect, embodiments of the disclosed subject matter can include a non-transitory computer readable medium. The non-transitory computer readable medium can include executable instructions operable to cause a data processing apparatus to receive one or more polygons associated with a region of interest, determine a plurality of sub-polygons that are contained within the one or more polygons, wherein each of the sub-polygons is associated with a unique code, and generate a first index system based on the tree structure, thereby providing an efficient mechanism to search whether a particular location is within the region of interest.

[0009] In any one of the embodiments disclosed herein, the unique code can include a location identifier based on a hierarchical encoding scheme, such as, for example, a geohash code.

[0010] In any one of the embodiments disclosed herein, the index can include one of a hash table or a probabilistic data structure.

[0011] In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for determining a difference between the index system and a previously-generated index system, and providing the difference to a computing device to update the previously-generated index system in the computing device.

[0012] In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for identifying a tree structure that models the unique code of the plurality of

sub-polygons based on a hierarchy of the unique codes and generating the first index system based on the tree structure.

**[0013]** In any one of the embodiments disclosed herein, the tree structure can include a branch node and a leaf node, and the branch node is associated with a geographic area that is larger than that of the leaf node, and wherein the apparatus, the method, or the non-transitory computer readable medium further includes modules, steps, or executable instructions for generating indices for the index system by traversing the tree structure from the branch node to the leaf node.

**[0014]** In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for determining a list of polygons that includes a particular sub-polygon, and associating the list of polygons with a leaf node of the tree structure corresponding to the particular sub-polygon.

**[0015]** In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for merging the tree structure, corresponding to the region of interest, with a second tree structure corresponding to a second region of interest, thereby providing a single index system modeling both the first tree structure and the second tree structure.

**[0016]** In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for removing one or more sub-polygons corresponding to a particular polygon from the index system.

**[0017]** In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for removing one or more sub-polygons from the plurality of sub-polygons to provide a reduced set of sub-polygons and generating the index system from the reduced set of sub-polygons, thereby reducing the size of the index system.

**[0018]** In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for receiving a second index system from another computing device and merging the first index system and the second index system by considering an overlap of sub-polygons corresponding to the first index system and the second index system.

[0019] In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for providing the index system to a computing device so that the computing device can use the index system to serve location queries.

[0020] In general, in an aspect, embodiments of the disclosed subject matter can include an apparatus, a method, and a non-transitory computer readable medium. The apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for receiving a location query from the client device, wherein the location query includes a location identifier associated with the client device, determining a query identifier corresponding to the location identifier, comparing the query identifier with the index system to determine that the location identifier provided by the client device is within the region of interest, and providing a service associated with the region of interest to the client device over the communication network.

[0021] In any one of the embodiments disclosed herein, the unique identifiers and the query identifier can include geohash codes.

[0022] In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for determining that the query identifier is represented in the index system and that the location identifier provided by the client device is within the region of interest

[0023] In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for comparing a first sequence of bits of the query identifier, corresponding to a lower-precision sub-polygon, to the index system before comparing a second sequence of bits of the query identifier, corresponding to a higher-precision sub-polygon.

[0024] In any one of the embodiments disclosed herein, wherein the index system comprises an index tree, and the apparatus, the method, or the non-transitory computer readable medium further includes modules, steps, or executable instructions for determining that the query identifier is within the region of interest when the first sequence of bits of the query identifier match a first index of the index system corresponding to a leaf node of the index system.

[0025] In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable

instructions for retrieving, from the index system, a polygon identifier associated with the query identifier, determining a group identifier associated with the polygon identifier, and providing the service associated with the campaign identifier to the client device over the communication network.

[0026] In any one of the embodiments disclosed herein, the apparatus, the method, or the non-transitory computer readable medium can include modules, steps, or executable instructions for retrieving, from the index system, a polygon identifier associated with the query identifier, and providing data associated with the polygon identifier to the client device over the communication network.

#### Description of the Figures

[0027] Various objects, features, and advantages of the present disclosure can be more fully appreciated with reference to the following detailed description when considered in connection with the following drawings, in which like reference numerals identify like elements. The following drawings are for the purpose of illustration only and are not intended to be limiting of the disclosed subject matter, the scope of which is set forth in the claims that follow.

[0028] FIG. 1 illustrates a diagram of a location query system in accordance with some embodiments.

[0029] FIGS. 2-3 illustrate how location identifiers can be used to represent a region of interest in accordance with some embodiments.

[0030] FIGS. 4A-4B illustrate properties of geohash codes and tiles in accordance with some embodiments.

[0031] FIG. 5 illustrates a tree of geohash codes in accordance with some embodiments.

[0032] FIGS. 6A-6D illustrate two collections of geohash tiles and the associated tree in accordance with some embodiments.

[0033] FIGS. 7A-7C illustrate a merging of the trees in accordance with some embodiments.

[0034] FIGS. 8A-8B illustrate a mechanism for generating the geohash tile index system in accordance with some embodiments.

[0035] FIG. 9 illustrates a process for serving location queries in accordance with some embodiments.

[0036] FIG. 10 illustrates a process for comparing a query geohash code to a geohash index system in accordance with some embodiments.

[0037] FIGS. 11A-11B illustrates a process for comparing a geohash code and a geohash index system in accordance with some embodiments.

[0038] FIG. 12 illustrates a process for building an index system associated with a set of regions of interest in accordance with some embodiments.

[0039] FIGS. 13A-13D illustrate a process of building an index tree corresponding to a region in accordance with some embodiments.

[0040] FIG. 14 illustrates a process for responding to a location query in accordance with some embodiments.

#### Detailed Descriptions

[0041] In the following description, numerous specific details are set forth regarding the systems and methods of the disclosed subject matter and the environment in which such systems and methods may operate, etc., in order to provide a thorough understanding of the disclosed subject matter. It will be apparent to one skilled in the art, however, that the disclosed subject matter may be practiced without such specific details, and that certain features, which are well known in the art, are not described in detail in order to avoid complication of the disclosed subject matter. In addition, it will be understood that the examples provided below are exemplary, and that it is contemplated that there are other systems and methods that are within the scope of the disclosed subject matter.

[0042] The disclosed apparatus, systems, and methods relate to a location query mechanism that can efficiently determine whether a target entity is located within a region of interest (ROI). At a high level, the location query mechanism can be configured to represent a ROI using one or more polygons. The location query mechanism can, in turn, divide (e.g., tessellate) the one or more polygons into sub-polygons. Subsequently, the location query mechanism can use the sub-polygons to build an index system that can efficiently determine whether a particular location is within any of the sub-polygons. Therefore, when a computing device queries whether a particular location is within the region of interest, the location query mechanism can use the index system to determine whether the particular location is within any of the sub-polygons.

[0043] In some embodiments, the disclosed location query mechanism can include three stages. The first stage includes representing the ROI with one or more polygons. The second

stage includes the generation of an index system for the one or more polygons. The index system generation process can involve receiving location descriptions of the one or more polygons and generating an efficiently query-able data structure for the location descriptions. The index system generation process can be performed off-line using a single computer or a cluster of computers. Therefore, the index system generation process may not interfere with an on-line (e.g., real-time) or high throughput (e.g. batch or real-time) operation of the location query response mechanism, as disclosed below.

**[0044]** The third stage includes a real-time query response mechanism for responding to location queries. For example, when the query response mechanism receives a location query, including a location identifier, from a client device, the query response mechanism can search the index system to determine whether the location identifier is associated with any of the polygons represented by the index system. If the location identifier is associated with one of the polygons, the query response mechanism can indicate the one or more polygons associated with the location identifier.

**[0045]** The disclosed location query mechanism is substantially more efficient compared to existing location query mechanisms. The disclosed location query mechanism can enable a server to serve queries in sub milliseconds and theoretically enable the processing of tens of thousands of queries per second per processing core.

**[0046]** The disclosed location query mechanism can be useful in the advertisement industry. For example, the advertisement display system on a mobile device can be configured to update and send the mobile device's geographic location information to an advertisement server. The advertisement server, in turn, can use the disclosed location query mechanism to identify relevant advertising campaigns related to the device's location, and use this information to serve highly contextual, location sensitive advertisement to the mobile device. Therefore, the query response mechanism can cause an advertisement to be sent to a target entity associated with the location identifier.

**[0047]** The disclosed location query mechanism can also be useful in mobile applications services. For example, sports-related mobile applications may provide a different interface on a user interface (e.g., a screen) of a mobile device depending on the device's location. For instance, a mobile device can update and send its geographic location information to the mobile application server. The mobile application server, in turn, can use the disclosed location query mechanism to identify relevant services related to the device's location, and use this information to serve highly contextual, location sensitive service to the mobile device.

[0048] The disclosed location query mechanism can also be useful in a variety of applications that processes data with location information. For example, processing data with location information from, for example, user logs, location tagged social network information (e.g. a stream of tweets), or similar data can benefit from additional contextual information such as whether the stated location is in a region of interest such as a shopping mall.

[0049] FIG. 1 illustrates a diagram of a location query system in accordance with some embodiments. The system 100 includes a host server 102, a communication network 104, and one or more client devices 106. The host server 102 can include a processor 108, a memory device 110, an index generation module 112, and a query response module 114. The host server 102 and the one or more client devices 106 can communicate via the communication network 104.

[0050] The communication network 104 can include the Internet, a cellular network, a telephone network, a computer network, a packet switching network, a line switching network, a local area network (LAN), a wide area network (WAN), a global area network, or any number of private networks currently referred to as an Intranet, and/or any other network or combination of networks that can accommodate data communication. Such networks may be implemented with any number of hardware and software components, transmission media and network protocols. Although FIG. 1 represents the network 104 as a single network, the network 104 can include multiple interconnected networks listed above.

[0051] A client 106 can include a desktop computer, a mobile computer, a tablet computer, a cellular device, or any other computing devices having a processor and memory. The client 106 can communicate with the host server 102 via the communication network 104. Although FIG. 1 represents the host server 102 as a single server, the host server 102 can include more than one server and can be part of a cloud-computing platform.

[0052] The processor 108 of the host server 102 can be implemented in hardware. The processor 108 can include an application specific integrated circuit (ASIC), programmable logic array (PLA), digital signal processor (DSP), field programmable gate array (FPGA), or any other integrated circuit. The processor 108 can also include one or more of any other applicable processors, such as a system-on-a-chip that combines one or more of a CPU, an application processor, and flash memory, or a reduced instruction set computing (RISC) processor. The memory device 110 of the processor 108 can include a computer readable medium, flash memory, a magnetic disk drive, an optical drive, a programmable read-only memory (PROM), and/or a read-only memory (ROM).

**[0053]** The index generation module 112 can be configured to generate an index system for one or more polygons. The index generation module 112 can maintain the generated index in the memory device 110 or provide the generated index to the query response module 114. The query response module 114 can be configured to respond to location queries in real-time. In some cases, the query response module 114 can reside in the host server 102. In other cases, the query response module 114 can reside in the client device 106. Also, the index generation module 112 and the query response module 114 need not reside on the same device.

**[0054]** In some embodiments, the index generation module 112 and/or the query response module 114 can be implemented in software stored in the memory device 110. The software stored in the memory device 110 can run on the processor 108 capable of executing computer instructions or computer code.

**[0055]** In some embodiments, the index generation module 112 and/or the query response module 114 can be implemented in hardware using an ASIC, PLA, DSP, FPGA, or any other integrated circuit. In some embodiments, the index generation module 112 and the query response module 114 can both be implemented on the same integrated circuit, such as ASIC, PLA, DSP, or FPGA, thereby forming a system on chip.

**[0056]** The index generation (“IG”) module 112 can be configured to use one or more location identifiers to represent a polygon. Location identifiers can be associated with any coordinate systems or hashing systems representing a region. More particularly, a polygon can be tessellated into a set of tiles (also referred to as sub-polygons). Each sub-polygon can cover geographic sub-region based on the desired level of precision, and can be associated with a location identifier. For example, a location identifier can include geohash code associated with a region of a predetermined precision or size. A region associated with a location identifier can be referred to as a tile or a sub-polygon. For example, a region associated with a geohash code can be referred to as a geohash tile or a geohash sub-polygon.

**[0057]** In some embodiments, location identifiers can be hierarchically organized. For example, certain types of location identifiers, such as geohash codes, can use a 32 subdivision system. Under the 32 subdivision system, a geohash code can be associated with a region that are covered by 32 other geohash codes, and each of the 32 other geohash codes can, in turn, be associated with a region that are covered by a plurality of other geohash codes. Therefore, the geohash codes can be represented as a tree. The hierarchy of the location identifiers can be determined based on a variety of factors, for example, a number of bits used

to represent a location identifier, a depth of the tree representing the hierarchy of location identifiers, and/or a breadth of the tree representing the hierarchy of location identifiers.

**[0058]** FIG. 2 illustrates how location identifiers can be used to represent a polygon in accordance with some embodiments. FIG. 2 includes a region of interest, which can include one or more polygons 200. A polygon can refer a geographic demarcation of an area. In some embodiments, the polygon can be represented graphically on a map; in other embodiments, the polygon can be represented by a plurality of subpolygons associated with location identifiers. For example, the polygon can be represented by a plurality of geohash tiles, including a first geohash tile 202 and a second geohash tile 204, each of which can be associated with a geohash code.

**[0059]** In some embodiments, a geohash tile can have one of several predetermined sizes. For example, the first geohash tile 202 is larger than the second geohash tile 204. As shown, smaller, higher precision, geohash tiles are used near the periphery of the polygon and larger, lower precision, geohash tiles are used in the interior of the polygon. The area of the polygon 200 is taken to be the collective area defined by all the tiles that form the polygon.

**[0060]** FIG. 3 also illustrates how location identifiers can be used to represent a polygon in accordance with some embodiments. FIG. 3 includes a region of interest, which is a polygon 300 having a circular shape surrounding a point of interest at the center of the polygon 300. As with FIG. 2, the polygon 300 is represented by a plurality of subpolygons associated with location identifiers. For example, the polygon 300 is represented by a plurality of geohash tiles, including a first geohash tile 302 and a second geohash tile 304, and each geohash tile can be associated with one of several predetermined sizes. As with FIG. 2, smaller, higher precision, geohash tiles are used to define the periphery of the polygon 300 and larger, lower resolution, geohash tiles are used in the interior of the polygon 300. Again, the area of the polygon 300 is taken to be the collective area defined by all geohash tiles that form the polygon 300.

**[0061]** In some embodiments, the IG module 112 can be configured to identify one or more sub-polygons (e.g., geohash tiles) that collectively represent a region comprising a plurality of polygons. The sub-polygon identification for a region can involve two steps. The first step of the sub-polygon identification can include receiving one or more polygons associated with a region. These polygons can be as simple as a coordinate (representing the location of a point-of-interest) enclosed by a circle of a given radius, or more complex shape,

like a multi-edged polygon representing a desired geographical area, or a point-of-interest, such as an airport.

**[0062]** The second step of the sub-polygon identification can include generating one or more sub-regions that are enclosed by the region (e.g., assembled polygons). For example, a geohash tile can be defined on a coordinate system, and can be considered a sub-region defined on that coordinate system. The IG module 112 can be configured to find sub-regions (e.g., geohash tiles), defined on the geohash coordinate system, that are entirely contained within the one of the assembled polygons. The IG module 112 can be configured to favor a region representation that uses larger sub-regions than smaller sub-regions so that a polygon can be represented with a small number of sub-regions.

**[0063]** The IG module 112 can be configured to identify such sub-polygons in an iterative manner. For example, as a first step, the IG module 112 can be configured to construct a set of geohash tiles having an identical, largest size such that this set of geohash tiles encompasses one or more polygons in the region of interest. Subsequently, the IG module 112 can be configured to test each geohash tile in the set of geohash tiles to determine if the particular geohash tile is completely within the associated polygon (e.g., without crossing the boundary of the associated polygon). If the particular geohash tile is completely within the polygon, the IG module 112 can keep the geohash tile. If the particular geohash tile is completely outside of the polygon, the IG module 112 can discard the particular geohash tile. If the particular geohash tile is partially within the polygon (e.g. crossing the boundary of the associated polygon,) the IG module 112 can break the particular geohash tile into a plurality of geohash sub-tiles.

**[0064]** Subsequently, the IG module 112 repeats the above process using the geohash sub-tiles. For example, the IG module 112 can determine, for each of the plurality of geohash sub-tiles, whether the sub-tile is completely within the associated polygon. If the sub-tile is completely within the polygon, the IG module 112 can keep the geohash sub-tile. If the sub-tile is completely outside of the polygon, the IG module 112 can discard the geohash sub-tile. If the sub-tile is partially within the polygon, the IG module 112 can further break the geohash sub-tile into a plurality of smaller tiles and repeat this process. The IG module 112 can perform this operation iteratively for each geohash tile to see if it ‘fits’ (i.e. does not intersect) the polygon, and if not, it recursively decreases the size of the geohash tile (e.g., increases the precision) to achieve a fit.

[0065] In some embodiments, there is a maximum level of precision defined for the geohash tiles (e.g., the smallest geohash tile that can be used to model the polygon), thereby providing a proper balance between the number of tiles and the level of fit for producing a good index. If a tile is reduced to the minimum size (e.g., the maximum precision) but still intersects the desired polygon, it is considered to be inside the polygon and included in the index.

[0066] In some embodiments, a location identifier associated with a sub-polygon can include a geohash code (e.g., a geohash code of the type defined in <http://geohash.org/>). A geohash code is a hierarchical spatial data structure that subdivides a region into tiles. A geohash code can include a sequence of bits that substantially uniquely identifies a location. In some cases, the sequence of bits can be encoded or can represent a sequence of characters. An example of a geohash code is a character sequence, “8z4fg.” In some embodiments, a set of geohash codes can exhibit hierarchical characteristics. For example, shorter geohash codes can be associated with a lower precision (e.g., shorter geohash codes are associated with larger geographic areas) whereas longer geohash codes can be associated with a higher precision (e.g., longer geohash codes are associated with smaller geographic areas). As a consequence of the gradual precision degradation based on the number of characters, nearby locations are often associated with similar prefixes. In some embodiments, the geohash codes that begin with the same characters can refer to the same geographic area. Two geohash codes that share a large number of prefix characters are associated with two locations that are in proximity.

[0067] FIGS. 4A-4B illustrate properties of geohash codes and tiles in accordance with some embodiments. FIG. 4A shows a geohash tile 402 corresponding to a geohash code “8z4fg.” The geohash tile 402 is centered on a geographical coordinate denoted by an X. X could represent any geographical coordinate, for example, latitude = 42 degrees North and longitude = 71 degrees West. FIG. 4B shows four geohash tiles 404A- 404D. Figure 4B also shows the geohash code associated with each tile (e.g., “8z4fgf” is the geohash code for tile 404A). The four geohash tiles 404A- 404D collectively cover the same area as the geohash tile 402. In fact, the four geohash tiles 404A- 404D represent four sub-divisions of the geohash tile 402. The X at the center of FIG. 4B represents the same geographical coordinate as the X shown in FIG. 4A.

[0068] Because the four geohash tiles 404A- 404D represent four sub-divisions of the geohash tile 402, the geohash codes for the four geohash tiles 404A- 404D can be longer than

the geohash code for the geohash tile 402, and the geohash codes for the four geohash tiles 404A- 404D can share characters with the geohash code for the geohash tile 402. For example, the geohash codes of the geohash tiles 404A- 404D can be each six characters long, whereas the geohash code the geohash tile 402 can be five characters long.

**[0069]** Furthermore, all geohash codes shown in FIG. 4B begin with the five character sequence “8z4fg”, which is identical to the geohash code of the geohash tile 402. This illustrates that geohash codes that begin with the same sequence relate to the same geographical area. Since the geohash codes of the geohash tiles 404A-404D begin with the same sequence as the geohash code of the geohash tile 402, the geohash codes of the geohash tiles 404A-404D are associated with tiles that fall within the geohash tile 402. Generally, if a geohash code corresponds to an original tile, then a concatenation of the geohash code with another character refers to a sub-tile within the original tile, as shown by the example of FIGS. 4A-4B.

**[0070]** In summary, the use of hierarchical encoding schemes for location identifiers, such as geohash codes, offer useful properties, such as an arbitrary precision (e.g., by adding as many characters as needed), a locality (e.g., similar prefixes are associated with nearby positions), and the ability to reduce precision (or increase an area coverage) by removing one or more characters from the end of the geohash code while still maintaining the spatial locality. These properties allow the disclosed indexing system to limit the amount of high precision indexing to only areas that require that level of detail.

**[0071]** In some embodiments, a collection of location identifiers, such as geohash codes (and their associated tiles) can be represented in a tree structure. FIG. 5 illustrates a tree of geohash codes in accordance with some embodiments. The tree of geohash codes (also referred to as a geohash tree) 500 relates to the geohash tiles illustrates in FIGS. 4A-4B. A geohash tree 500 can include branch nodes, such as a node 502, and leaf nodes, such as a node 504. A node is a branch node if it connects to a lower node (either a lower branch node or a leaf node). For example, the branch node 502 connects to another branch node 506. A node is a leaf node if it does not connect to any lower nodes. For example, the node 504 is a leaf node because it does not connect to any lower nodes. Some of the nodes in FIG. 5 can be associated with a geohash tile in FIG. 4. For example, the node 508 can be associated with the geohash tile in FIG. 4A. As another example, the nodes 510, 512, 514, and 504 can be associated with the geohash tiles 404A, 404B, 404C, and 404D, respectively.

[0072] FIGS. 6A-6D illustrate two collections of geohash tiles and the associated tree in accordance with some embodiments. FIG. 6A shows a single “large” geohash tile 602, whereas FIG. 6B shows a set of smaller geohash tiles, including tiles 604, 606. In FIG. 6B, the dashed periphery is the same size as that of the large tile 602 shown in FIG. 6A. FIG. 6B represents a case in which the large tile 602 does not fit a polygon well and so was split up into smaller tiles, including 604, 606, as previously discussed with respect to FIGS. 2-3. The single large tile 602 can be represented as a tree structure that includes only one single leaf node 608, as illustrated in FIG. 6C. On the other hand, the set of tiles shown in FIG. 6B can be represented as a tree structure that includes both branch node 608 and leaf nodes 610, 612, 614, 616, as illustrated in FIG. 6D.

[0073] Once the IG module 112 defines a polygon that describes a region of interest associated with an advertising campaign, the IG module 112 can generate one or more tree structures that describe the collection of geohash tiles in the polygon. The IG module 112 can repeat this process for each set of geohash tiles in each polygon of interest.

[0074] In some embodiments, the IG module 112 can merge multiple geohash trees so that multiple geohash trees can be represented using a compact representation. This feature can be useful when two or more computers are configured to generate multiple geohash trees in a distributed manner, for example, simultaneously. At a high level, when two polygons have intersecting geohash tiles, a lower precision geohash tile that encompasses other higher precision geohash tiles can be marked as a leaf, while the higher precision geohash tiles are discarded from the geohash tree. The net result is an optimized set of different precision geohash tiles that can be used to represent the set of disjoint polygons that make up a specified geo-targeted advertising campaign. This merging operation does not necessarily lose precision information because, if regions not covered by the higher precision geohash tiles are within the polygon according to a first geohash tree, there is no need to differentiate regions in the higher precision geohash tiles and regions outside of the higher precision geohash tiles.

[0075] FIGS. 7A-7C illustrate a merging of the trees in accordance with some embodiments. FIGS. 7A-7B illustrate two trees. The first tree 702 corresponds to a first polygon in the region of interest; the second tree 706 corresponds to a second polygon in the region of interest. The first tree 702 includes a geohash tile “4fg” 704, and the second tree 706 includes subtiles of the geohash tile 704. Therefore, the region of interest includes both the geohash tile “4fg” 704, as indicated by the first tree 702, and the subtiles of the geohash

tile 704, as indicated by the second tree 706. The IG module 112 can merge these two trees to generate a third tree 708 in which the subtiles of the geohash tile 704 is subsumed by the geohash node 704. Therefore, the third tree 708 can represent the region of interest having both the first polygon and the second polygon with a less number of nodes in the tree compared to the combined number of nodes in the first tree 702 and the second tree 706.

**[0076]** Subsequently, the IG module 112 can use the merged geohash tree to generate an index system for the geohash tiles. The IG module 112 can generate the index system by walking down the geohash tree from the top branch node to the leaf nodes in hierarchical order.

**[0077]** In some embodiments, the IG module 112 can reduce the resolution of a geohash tree to reduce the size of the geohash tree and to increase the query speed of the index system associated with the geohash tree. For example, referring to FIG. 7B, the IG module 112 can truncate, in the geohash tree 706, all nodes below the “g” node 704. This way, the IG module 112 can reduce the size of the geohash tree 706 at the expense of reducing the resolution of the polygon represented by the geohash tree 706.

**[0078]** FIGS. 8A-8B illustrate a mechanism for generating the geohash tile index system in accordance with some embodiments. FIG. 8A illustrates a geohash tree 800 associated with three geohash tiles: a first geohash tile 802 with a geohash code “9q5d”, a second geohash tile 804 with a geohash code “9q5dt”, and a third geohash tile 806 with a geohash code “9q5dw.” The IG module 112 can walk down the geohash tree 800 recursively to generate the geohash index system, as illustrated in FIG. 8B. The IG module 112 can start from the top (e.g., the root node) of the tree 800, and, at each stage (e.g., each level of the tree or a geodesic distance level from the root node), emit an index corresponding to the value of the node at that stage.

**[0079]** For example, in the first step, the IG module 112 can generate “9” as an index, since “9” is the value of the root node of the tree 800. Then the IG module 112 can walk down to the next node (e.g., the next level) and generate “q” as an index. Then the IG module 112 can walk down to the next node and generate “5” as an index. Then the IG module 112 can walk down to the next node and generate “d” as an index. Then the IG module 112 can walk down to the next node and generate “t” and “w” as indices associated with that level. In some embodiments, the IG module 112 can walk the tree in a depth-first-search manner; in other embodiments, the IG module 112 can walk the tree in a breadth-first-search manner.

**[0080]** FIG. 8B illustrates the generated geohash index system based on the geohash tree of FIG. 8A. In some embodiments, the geohash index system can be represented using a flat hierarchical data representation. In some cases, the flat hierarchical data representation can include a jump table. The jump table is, in effect, an unrolled representation (e.g., rasterized representation) of the geohash tree 800. The unrolled representation can be a single-layer data structure that can expedite the search for relevant information. In other cases, the flat hierarchical data representation can include a skip list.

**[0081]** In some embodiments, the geohash index system can be represented as a plurality of data structure nodes 808-816. Each data structure node can correspond to a set of nodes at the same level (e.g., the same distance from the root node) in the corresponding tree. For example, all nodes in level 3 (e.g., the number of shortest-path edges between the root node and a candidate node is 3), can be represented as a data structure node 814 having three values: the length of the geohash nodes in the tree at the current level 818, the value(s) 822 of the geohash nodes in the tree at the current level 818, and the number of “jumps” 820 to be performed to reach the data structure corresponding to the geohash nodes in the tree at the next level.

**[0082]** An alternative approach to encoding the data is to populate a probabilistic data structure, such as a bloom filter. Both approaches have merit and offer different tradeoffs. The flattened index approach described above has the characteristics of giving a more deterministic answer to the question, but at a potentially increased memory footprint when compared to the bloom filter approach. The bloom filter, on the other hand, is potentially more compact and would not require merging tree structures but has a probabilistic margin of error and thus can return false positives and also has a greater impact of the processor’s ability to prefetch memory pages.

**[0083]** An alternative approach to encoding the data is to use a hash table.

**[0084]** Once the IG module 112 generates the geohash index system, the IG module 112 can store the geohash index system in the memory device 110. Subsequently, the query response (“QR”) module 114 can use the stored geohash index system to serve location queries from clients 106.

**[0085]** FIG. 9 illustrates a process used by the QR module 114 for serving location queries in accordance with some embodiments. In step 902, the QR module 114 can be configured to receive a location query from a client 106, requesting the QR module 114 to

determine whether the client 106 is within a polygon. To this end, the location query can include the location identifier of the client 106, indicating a location of the client 106. The location identifier can take the form of a coordinate, for example, (longitude, latitude) pair. The client 106 can determine its location identifier based on a location determination mechanism. The location determination mechanism can include a Global Positioning System (GPS) technique, a cellular tower triangulation technique, an Internet Protocol (IP) address-based location determination technique, and/or any other suitable techniques for determining the location of the client 106.

**[0086]** In step 904, the QR module 114 can be configured to convert the location identifier into a geohash code. In some embodiments, the QR module 114 can be configured to generate the highest-precision query geohash code corresponding to the location identifier (e.g., a geohash code that most precisely identifies a location associated with the location identifier.) In some embodiments, the precision of the query geohash code can be higher than the maximum precision level of geohash codes summarized by the geohash index system.

**[0087]** In step 906, the QR module 114 can compare the query geohash code to the geohash index system, and in step 908, the QR module 114 can determine, based on the comparison, whether the received location identifier is within a polygon modeled by the geohash index system. If the received location identifier is within the polygon, then in step 910, the QR module 114 can provide the identifier of the polygon that matched to the location identifier (e.g., so as to match the received location identifier with a particular advertising campaign). If the received location identifier is not within the polygon, then in step 912, the QR module 114 can indicate that the received location identifier does not correspond to the polygon. In some cases, in step 910, the host server 102 can cause an advertisement associated with that polygon to be sent to the client 106. In some embodiments, this system allows such processing for multiple location identifier received from multiple clients 106 to be performed quickly to determine whether the location identifiers provided by the client devices 106 are in any of the polygons.

**[0088]** In some embodiments, in step 908, the QR module 114 can determine whether the received location identifier is within an polygon modeled by the geohash index system by comparing the query geohash code of the location identifier to the geohash index system. In some cases, the QR module 114 can be configured to compare characters corresponding to larger geohash tiles (e.g., lower-precision geohash tiles) before comparing characters

corresponding to smaller geohash tiles (e.g., higher-precision geohash tiles.) For example, the QR module 114 can retrieve the first character of the query geohash code and compare the first character to the root node (e.g., the highest node) in the geohash index system, modeled by the jump table and offsets. If the first character of the query geohash code matches one of the root nodes in the geohash index system, the QR module 114 can determine if the one of the root nodes represents a leaf node. If so, the QR module 114 can indicate that the query geohash code is associated with a polygon modeled by the index system and move to step 910. If the one of the root nodes does not represent a leaf node, the QR module 114 can move to the next character (e.g., a character adjacent to the first character), and compare the new character with values in one or more nodes coupled to the one of the root nodes (e.g., one or more children of the one of the root nodes.)

**[0089]** This process is iterated until (1) the QR module 114 does not find a match between a character and a value of the nodes in the level (e.g., the depth level of the geohash tree) corresponding to the character, or (2) the QR module 114 reaches the leaf node. If, at any point in walking down the geohash tree, the QR module 114 reaches a node where the character of the query geohash code does not match the values in the geohash index system, then the QR module 114 can declare a no-match and proceed to step 912. If the QR module 114 reaches the leaf node and the value of the leaf node matches a corresponding character in the query geohash code, then the QR module 114 can indicate a match between the query geohash code and the geohash index system, and proceed to step 910. If the QR module 114 reaches the leaf node and the value of the leaf node does not match a corresponding character in the query geohash code, then the QR module 114 can indicate a no-match between the query geohash code and the geohash index system and proceed to step 912.

**[0090]** FIG. 10 illustrates a process for comparing a query geohash code to a geohash index system in accordance with some embodiments. In this example, the QR module 114 can communicate with the geohash index system 1002 to determine whether a query geohash code “9q5f” 1004 is within a polygon modeled by the geohash index system 1002. In the first step, the QR module 114 can take the first character “9” of the query geohash code 1004 and compare it against the root data structure 808 of the geohash index system 1002. Since the value of the first character “9” matches the value of the root data structure 808, the QR module 114 can move onto the second data structure 810. In the second step, the QR module 114 can take the second character “q” of the query geohash code 1004 and compare against the second data structure 810 of the geohash index system 1002. Since the value of the

second character “q” matches the value of the second data structure 810, the QR module 114 can move onto the third data structure 812.

**[0091]** In the third step, the QR module 114 can take the third character “6” of the query geohash code 1004 and compare against the third data structure 812 of the geohash index system 1002. Since the value of the third character “6” does not match the value of the third data structure 812, the QR module 114 can determine that the query geohash code 1004 is not within the polygon modeled by the geohash index system 1002.

**[0092]** If, instead, the query geohash code is “9q5dt”, then the QR module 114 will find a match at each data structure node in the geohash index system 1002, and therefore, the QR module 114 would indicate that the location corresponding to the query geohash code is “9q5dt” is within the polygon modeled by the geohash index system 1002.

**[0093]** In some cases, a single comparison between the query geohash code and the geohash tree can be sufficient reveal whether the receive location identifier is within a polygon. For example, if the query geohash code of the location identifier lies within a tree having only a single leaf node, then the comparison of the first character of the query geohash code to the tree structure can reveal a hit and it will be known that the client 106 is located within an area of interest. FIG. 11A illustrates a scenario in which a single comparison can reveal a hit in accordance with some embodiments. In FIG. 11A, the rectangle 1102 corresponds to a single geohash tile, which is represented as a tree structure containing only a single leaf, as illustrated in FIG. 6C. The X in FIG. 11A represents the location of the client 106 from which the location identifier is received. Due to the nature of query geohash codes, the first character of the client’s geohash code is identical to the first character of the geohash code corresponding to the geohash tile 1102. Therefore, a comparison of the first character of the client’s query geohash code to the top level of the geohash tree can reveal a match.

**[0094]** In some cases, multiple comparisons between the query geohash code and the geohash tree may be needed to reveal whether the receive location identifier is within a polygon. FIG. 11B illustrates a scenario in which multiple comparisons may be needed to reveal a hit in accordance with some embodiments. In FIG. 6B, the set of geohash tiles is not represented by a single leaf. Rather, the geohash tree for that set of tiles includes multiple nodes, some of which are branches and others of which are leaves, as illustrated in FIG. 6D. Therefore, the process of determining whether the receive location identifier is within an polygon includes comparing the query geohash code of the location identifier to the geohash tree, one character at a time, from the top node of the geohash tree to the bottom node of the

geohash tree. If the first (e.g., the most significant) character of the query geohash code matches one of the geohash trees in the index system, then the query geohash code is compared against that tree structure's lower nodes, one node at a time.

**[0095]** Once the QR module 114 identifies at least one character in the query geohash code that is not represented by the geohash index system, then the QR module 114 can move to step 914, indicating that the location identifier is not within an polygon. Any mismatch means the phone does not lie within the area of interest. On the other hand, if the QR module 114 reaches the leaf node and the character in the leaf node matches the corresponding character in the query geohash code, then the QR module 114 can indicate that the location identifier is within the polygon modeled by the geohash tree.

**[0096]** In some embodiments, the index system (e.g., the geohash tree) is designed to yield fast comparison performance, returning an answer in microseconds using a single core of a commodity server. In some embodiments, the index system can be designed to be re-entrant, so lookups can scale out and take advantage of all available cores in the system without any adverse performance impacts due to lock conflicts.

**[0097]** In some embodiments, the index system can include information on polygons corresponding to each geohash tile modeled by the index system. For example, the IG module 112 can determine a list of all polygons that includes a particular geohash tile, and associate that list to the leaf node corresponding to the particular geohash tile. This way, the index system can maintain a correspondence between a geohash tile and all polygons that includes the geohash tile. Subsequently, when the QR module 114 finds a match between a query geohash code and the index system, the QR module 114 can return not only an identifier associated with the polygon, but also the specific polygons within the polygon that contributed to the match between the query geohash code and the index system.

**[0098]** In some embodiments, a polygon can include one or more group identifiers. Group identifiers, can, for example relate to a campaign for advertisements. For example, a campaign can include a plurality of polygons that collectively define areas to which a particular advertisement campaign can be targeted. Each campaign can be associated with an identifier that identifies the associated advertisement. In some embodiments, an index system corresponding to a campaign can be stored in a single file. Each campaign file may contain many thousands of individual polygons (also referred to as geofences). For example, a campaign file can include all McDonalds' locations within a 10 km radius of New York Penn Station.

**[0099]** In some embodiments, the IG module 112 can be configured to merge multiple indices into a single index. In some embodiments, the IG module 112 can be configured to remove a group from an existing index system. For example, the IG module 112 can traverse the existing index system (e.g., the tree), remove geohash codes associated with the group to be removed, and recursively rebuild the portion of the existing index system (e.g., a sub-tree) with the remaining geohash codes.

**[0100]** In some embodiments, the IG module 112 can provide the geohash index system to the client 106 so that the client 106 can directly serve location queries from other devices, such as mobile devices. In some cases, the IG module 112 can use a delta compression technique to provide only modified parts of the geohash index system to the client 106. The client 106 can use double buffering techniques to change its current index system to bring it up to date with the new geohash index system, without impacting existing query performance. For example, the client 106 can maintain the current index system in memory and then load the new index system into memory while still processing requests by based on the current index system. Once the new index system is fully loaded into memory and ready to respond to queries, a pointer to the current index system can instead reference the new index system such the requests are processed by referncing the new index system. Once that step is complete, the current index system can be deallocated from memory. To accommodate these processes, it may be desirable for the client 106 to have enough memory to store both the current index system and the new index system.

**[0101]** In some embodiments, the host server 102 can provide an application programming interface (API) or web interface to allow advertisement service entities to create a geohash index system for advertisement. For example, the API or web interface can allow the entities to generate a group using simple query criteria such as ‘within a 1 mile radius of businesses of type x in region y’. Once the entities select one or more polygons that represent the advertisement group, the host server 102 can use the above described method to generate the index system for the group. The index system can be represented as an index file. Then the host server can transmit the index file to clients 106 (e.g., group (e.g. campaign) owner / advertisement networks machines / servers) with a QR module 114 so that the index file can be incorporated into the existing index system in the client 106. This allows the QR module 114 in the client devices 106 to directly serve the location queries from mobile devices, instead of requesting the host server 102 to resolve the location queries.

[0102] In some embodiments, the IG module 112 can be configured to generate an index system that is capable of returning a set of polygons intersecting with a particular location. Such an index system can be useful in a variety of applications. In particular, such an index system can facilitate a mechanism for providing information and grouping identifiers of a particular location, collectively referred to as a payload data of a particular location. For instance, the QR module 114 can use such an index system to determine one or more identifiers of polygons associated with the particular location and to use the one or more identifiers of polygons to retrieve the payload data for the particular location from a database.

[0103] More particularly, a grouping identifier can be associated with a group of polygons to be represented together. The grouping identifier can be useful to associate the same property to each of the polygons identified in the group of polygons. For example, an advertising campaign by Carl's Jr can target its advertisement to users who are near either a McDonald's restaurant or a Burger King restaurant. In this case, the advertising campaign can be tagged with a campaign identifier "carls"; one or more polygons associated with the McDonald's restaurant can be tagged with a grouping identifier "mcd"; and one or more polygons associated with the Burger King restaurant can be tagged with a grouping identifier "bk". Furthermore, the grouping identifiers "mcd" and "bk" can be associated with the advertising campaign identifier "carls." Subsequently, when an advertiser wants to post advertisements for Carl's, the advertiser can use the association between the campaign identifier "carls" and the grouping identifiers "mcd" and "bk" to identify all polygons to be associated with the advertisement campaign.

[0104] FIG. 12 illustrates a process for building an index system that is capable of returning a set of polygons intersecting with a particular location in accordance with some embodiments. The index system can represent one or more regions of interest, where each region of interest can include one or more polygons, and each polygon can be associated with a unique polygon identifier. The polygon identifier can include a string of characters, e.g., including numbers, which refers to a specific polygon. The polygon identifier can be a UUID having, for example, 128 bits, or an integer having, for example, 32 or 64 or some other number of bits. The polygon identifier can point to a specific entity of interest, for example, a specific McDonald's restaurant at a specific address rather than McDonald's restaurants in general, which may be represented by an index or a group identifier.

[0105] When the IG module 112 is configured to generate an index system for one region of interest, the IG module 112 is configured to generate an index system for that region by

iterating steps 1202-1206. When the IG module 112 is configured to generate an index system for more than one region of interest, the IG module 112 is configured to generate an index system for each region independently by iterating steps 1202-1206, and consolidate the index systems for each region as a post-processing step in step 1208.

**[0106]** In some embodiments, the index system for a region can have a tree structure. Therefore, the index system for a region can be referred to as an index tree. A node in the index tree can be associated with a sub-region of a region. Each node in the index tree can also be associated with one or more identifiers of polygon(s) that intersect with the sub-region associated with the node.

**[0107]** In some embodiments, the IG module 112 is configured to process the index tree so that one or more polygon identifiers associated with a node can be represented succinctly. For example, the IG module 112 has a mechanism for declaring a particular polygon identifier as a leaf identifier. When the IG module 112 declares a particular polygon identifier as a leaf identifier at a particular node, then all children nodes of the particular node are deemed to be associated with the particular polygon identifier. This way, the IG module 112 obviates the need to explicitly associate the leaf identifier with every child node, thereby reducing redundant associations of polygon identifiers in the index tree.

**[0108]** More particularly, in step 1202, the IG module 112 is configured to represent (e.g., tessellate) a polygon in a region using a plurality of tiles (e.g., sub-polygons). A sub-polygon is designed to cover an area whose size depends on a predetermined level of precision associated with the sub-polygon. For example, when the precision is low, the sub-polygon covers a large area; when the precision is high, the sub-polygon covers a small area.

**[0109]** The IG module 112 is also configured to associate a sub-polygon with an identifier of the polygon from which the sub-polygon is derived. For example, when a polygon is divided into 32 sub-polygons, each of the sub-polygons is associated with a polygon identifier of the original polygon. If a region includes more than one polygon, this process is repeated for each polygon in the region. Therefore, a single sub-polygon can be associated with a plurality of polygon identifiers. Subsequently, the sub-polygons used to represent polygons in the region can be grouped together to represent the region.

**[0110]** In step 1204, the IG module 112 is configured to recursively subdivide the region into sub-regions and associate each sub-region with a polygon identifier covering the sub-region. More particularly, the IG module 112 is first configured to identify one or more

unique identifiers of polygons in the region. Then, the IG module 112 is configured to iteratively subdivide the region into sub-regions (and sub-regions into smaller sub-regions), and assign one or more of the unique polygon identifiers to a sub-region if the sub-region intersects with polygons represented by the unique polygon identifiers. Since each sub-polygon is associated with one or more polygon identifiers, the IG module 112 can determine a set of polygon identifiers associated with each sub-region.

**[0111]** As the IG module 112 iteratively subdivide a region into sub-regions (and sub-regions into smaller sub-regions), the IG module 112 can build an index tree corresponding to the region. FIGS. 13A-13D illustrate a process of building an index tree corresponding to a region in accordance with some embodiments.

**[0112]** FIG. 13A illustrates a region 1300 that includes a plurality of tiles (e.g., sub-polygons, illustrated as small squares), where each tile is associated with a location identifier, such as a geohash code. The region 1300 can represent one or more polygons that are illustrated using dotted tiles. For example, a dotted tile 1304 indicates a tile that belongs to one or more polygons in the region 1300; an empty tile 1306 indicates a tile that is outside of one or more polygons in the region 1300. For the purpose of illustration, the following exemplary descriptions deem the region 1300 to include a single polygon 1302, but the descriptions are also applicable to cases in which the region 1300 includes multiple polygons.

**[0113]** To generate an index tree for the region 1300, the IG module 112 can instantiate an index tree 1308 with a single root node corresponding to the entire region 1300. FIG. 13B illustrates the region 1300 and the index tree 1308 having a single root node 1310 corresponding to the entire region 1300. Subsequently, the IG module 112 can subdivide the region 1300 into two sub-regions 1312, 1314, and add nodes 1316, 1318 corresponding to the sub-regions 1312, 1314. Also, the IG module 112 can associate the nodes 1316, 1318 with polygon identifiers of polygons that intersect with the sub-regions 1312, 1314, respectively.

**[0114]** Subsequently, the IG module 112 can determine whether a polygon covers an entire area represented by a sub-region 1312 or 1314. If so, the IG module 112 can mark the polygon identifier of that polygon as a leaf identifier for that sub-region (or a node corresponding to that sub-region), which indicates that all sub-trees rooted at that sub-region include that leaf polygon identifier. In the example shown in FIG. 13B, every tile in the sub-region 1312 includes the same polygon identifier corresponding to the polygon 1302. Therefore, the polygon 1302 covers an entire area represented by the sub-region 1312. Thus,

the polygon identifier corresponding to the polygon 1302 in the node 1316 is a leaf identifier at the node 1316.

**[0115]** In some embodiments, if all polygon identifiers associated with a node is a leaf identifier, then the IG module 112 can stop building the index tree 1308 for the sub-region corresponding to that node (e.g., stop sub-dividing the sub-region corresponding to that node). In the example shown in FIG. 13B, since the region 1300 includes only a single polygon 1302, all polygon identifiers associated with the node 1316 is a leaf identifier. Therefore, the IG module 112 can stop sub-dividing the sub-region 1312 corresponding to the node 1316. On the other hand, not every tile in the sub-region 1314 includes the same polygon identifier corresponding to the polygon 1302. For instance, some of the tiles in the sub-region 1314 does not intersect with the polygon 1302. Therefore, the polygon identifier associated with the node 1318 is not a leaf identifier. Thus, the IG module 112 is configured to further sub-divide the sub-region 1314.

**[0116]** FIG. 13C illustrates a division of the sub-region 1314 into sub-regions 1320, 1322, 1324, and 1326. Because only the sub-regions 1324, 1326 include a tile of the polygon 1302, the IG module 112 is configured to generate only two additional nodes 1328, 1330 in the index tree 1308 associated with the sub-regions 1324, 1326, respectively, and discard sub-regions 1320, 1322 from further processing.

**[0117]** Furthermore, the IG module 112 can stop sub-dividing the sub-region 1324 corresponding to the node 1328. Every tile in the sub-region 1324 includes the polygon identifier of the polygon 1302. Therefore, the polygon identifier of the polygon 1302 in the node 1328 is a leaf identifier. Furthermore, since the region 1300 includes only a single polygon 1302, all polygon identifiers associated with the node 1328 is a leaf identifier. Therefore, the IG module 112 can stop sub-dividing the sub-region 1324 corresponding to the node 1328.

**[0118]** On the other hand, the IG module 112 is configured to further sub-divide the sub-region 1326. Not every tile in the sub-region 1326 includes the polygon identifier of the polygon 1302. For example, some of the tiles in the sub-region 1326 does not intersect with the polygon 1302. Therefore, at least one of the polygon identifiers in the node 1330 is not a leaf polygon identifier. Therefore, the IG module 112 is configured to further sub-divide the sub-region 1326.

**[0119]** FIG. 13D illustrates a division of the sub-region 1326 into sub-regions 1332, 1334, 1336, and 1338. Through a similar process as outlined above, the IG module 112 is configured to add two additional nodes 1340, 1342 to the index tree 1308, corresponding to sub-regions 1336, 1338, respectively.

**[0120]** In some embodiments, once the IG module 112 completes the index generation process, the IG module 112 can be configured to traverse the index tree 1308 from the leaf nodes (e.g., nodes 1340, 1342) to the root node (e.g., node 1310) to reduce the number of polygon identifiers associated with the index tree 1308. The IG module 112 is configured to determine whether all children nodes of a particular node (also referred to as a parent node) share the same polygon identifier. If so, the IG module 112 is configured to remove that polygon identifier from all children nodes, associate the parent node with that polygon identifier, and declare that polygon identifier as the leaf identifier at the parent node. This reduction process can reduce the number of polygon identifiers at the highest precision nodes (e.g., nodes furthest away from the root node).

**[0121]** In some embodiments, the IG module 112 can reduce a number of bits used to represent a polygon identifier in the index tree. To this end, the IG module 112 is configured to present a polygon identifier in a child node as an offset into a set of polygon identifiers in the parent node. For example, suppose that a parent node is associated with three polygon identifiers: [021y4bcfjpk26rsx, pr2swz25xyqebc13, fm0qrx36zmn79fjppq], and has a child node that is associated with two polygon identifiers: [021y4bcfjpk26rsx, fm0qrx36zmn79fjppq]. Instead of actually writing out the polygon identifiers in the child node, the IG module 112 can be configured to represent the polygon identifiers in the child node as an index into the three polygon identifiers in the parent node. Under this scheme, the IG module 112 can represent the two polygon identifiers in the child node as [1, 3]. This representation can reduce the number of bits used to represent polygon identifiers in the index tree.

**[0122]** Once the index tree is constructed, the IG module 112 can encode the index tree (e.g., the polygon identifiers in each node of the index tree, the set of leaf polygon identifiers) into an index system, as illustrated in FIG. 8B in accordance with some embodiments.

**[0123]** Once the IG module 112 completes the index system generation for a region, the IG module 112 can encode all the payload data associated with the region so that the data associated with the region (or a polygon within the region) can be retrieved quickly. To limit the magnitude of the offsets / jump indexes encoded at each sub-region level (e.g., each level

in the index tree), the IG module 112 is configured to encode each level's data into a separate substream, including a sub-index or an independent portion of the index.

[0124] Subsequently, the IG module 112 can write out this substream data (for each level in the index tree) at the head of the total payload data representing the entire region represented by the index tree. The IG module 112 can write out all polygon identifiers represented within the given region. Then, the IG module 112 can encode the hierarchical region/sub-region index tree, using a technique such as a geohash encoding. These operations can complete the index generation process and the payload database generation process for a region.

[0125] If there are more than one regions of interest, the IG module 112 can repeat steps 1202 – 1206 for each region of interest, and generate an index system for each region. Once the index system for each region is constructed, in step 1208, the IG module 112 is configured to merge the index systems into a single master index system so that the single master index system can represent all regions of interest. In some embodiments, the single master index system can have a tree structure, and the tree structure can be based on geohash codes. Each leaf node of the single master index system can correspond to a region-level index system generated in steps 1202-1206.

[0126] The index generation process of FIG. 12 is beneficial because region-level index systems can be generated in parallel using a distributed computing system. Since regions can be independent of one another, the IG module 112 can assign a single computer to generate an index system for one region without worrying about an index system for another region. This allows for a parallel computation of region-level index systems, thereby providing a computationally scalable mechanism for generating indices for many regions of interest. Furthermore, the index generation process of FIG. 12 allows the IG module 112 to quickly update the index system by building and encoding the index tree of only the region that needs to be updated.

[0127] The QR module 114 can use the master index system generated by the IG module 112, as illustrated in FIG. 12, to respond to a location query from a client. In particular, the QR module 114 can use the master index system to determine one or more polygons intersecting a location identified in the location query and provide any payload data associated with the one or more polygons to the client.

**[0128]** FIG. 14 illustrates a process 1400 for responding to a location query in accordance with some embodiments. In step 1402, the QR module 114 receives the location query and extracts a location identifier, such as a [latitude, longitude] pair or a geohash code, from the location query. The location associated with the location identifier is referred to as a target location. Then the QR module 114 can identify the region that includes the target location, and retrieve, from the master index system, the region-level index system corresponding to that region. Subsequently, the QR module 114 can allocate memory space to hold the set of polygon identifiers represented by the region, as encoded in the index tree corresponding to the region. The memory space can be used to hold identifier(s) of polygon(s) that intersect with the target location.

**[0129]** In step 1404, the QR module 114 can walk down the region-level index system to find sub-regions of the region that also intersect with the target location. As the QR module 114 traverses down the index tree, at each node during the traversal, the QR module 114 can collect leaf polygon identifiers associated with the node. Then the QR module 114 can narrow the potential set of potential leaf identifiers that might be found in subsequent iterations. For example, as the QR module 114 walks down the node hierarchy of the index tree, the set of polygon identifiers relevant to a particular node corresponding to a particular sub-region is restricted to the polygon identifiers associated with the parent node of the particular node. Once the QR module 114 reaches the leaf node of the region-level index system (e.g., the highest precision level of the region-level index system), the QR module 114 can terminate the traversal of the region-level index system. The resulting set of leaf polygon identifiers represents the set of polygons intersecting the location identifier.

**[0130]** In step 1406, once the QR module 114 identifies all polygons intersecting the target location, the QR module 114 can retrieve payload data associated with the polygons using their identifiers. For example, the QR module 114 can request a database table or a hash table to retrieve any data, such as the name or address, associated with a particular polygon identifier. Subsequently, the QR module 114 can provide, to the client that sent the location query, the set of polygon identifiers associated with the target location and any retrieved data associated with the polygons.

**[0131]** While the foregoing embodiments have been illustrated primarily using geohash codes and geohash tiles, the foregoing embodiments can use other location identification mechanisms as well.

**[0132]** For example, in any of the embodiments disclosed herein, a region can be represented by any type of a location identifier associated with a hierarchical location identifier system, including, for instance, a hash-based location identifier and/or a quad tree-based location identifier. Under the hierarchical location identifier systems, a concatenation of a location identifier, associated with a region, with one additional bit (or character) can refer to one of a predetermined number of sub-regions contained within the region. For instance, a concatenation of a location identifier, associated with a region, with one additional bit can refer to one of 4 sub-regions, 8 sub-regions, or 16 sub-regions contained within the region. Also, for instance, a concatenation of a location identifier, associated with a region, with one additional sequence of bits or one additional character can refer to one of 4 sub-regions, 8 sub-regions, or 16 sub-regions contained within the region.

**[0133]** As another example, in any of the embodiments disclosed herein, a region can be represented by any type of a location identifier that can be associated with a particular geographic/physical precision. For instance, a data structure, such as a probabilistic data structure including a bloom filter, may be associated with one of a predetermined set of precisions (e.g. 1m, 3m, 5m, and 10m) and can encode a location identifier associated with the one of the predetermined set of precisions in, for example, as few bits as possible.

**[0134]** In some embodiments, the client 106 can include user equipment of a cellular network. The user equipment communicates with one or more radio access networks and with wired communication networks. The user equipment can be a cellular phone having phonetic communication capabilities. The user equipment can also be a smart phone providing services such as word processing, web browsing, gaming, e-book capabilities, an operating system, and a full keyboard. The user equipment can also be a tablet computer providing network access and most of the services provided by a smart phone. The user equipment operates using an operating system such as Symbian OS, iPhone OS, RIM's Blackberry, Windows Mobile, Linux, HP WebOS, and Android. The screen might be a touch screen that is used to input data to the mobile device, in which case the screen can be used instead of the full keyboard. The user equipment can also keep global positioning coordinates, profile information, or other location information.

**[0135]** The client 106 also includes any platforms capable of computations and communication. Non-limiting examples can include computers, netbooks, laptops, servers, and any equipment with computation capabilities. The client 106 is configured with one or more processors that process instructions and run software that may be stored in memory.

The processor also communicates with the memory and interfaces to communicate with other devices. The processor can be any applicable processor such as a system-on-a-chip that combines a CPU, an application processor, and flash memory. The client 106 can also provide a variety of user interfaces such as a keyboard, a touch screen, a trackball, a touch pad, and/or a mouse. The client 106 may also include speakers and a display device in some embodiments.

**[0136]** In some embodiments, the host server 102 can reside in a data center and form a node in a cloud computing infrastructure. The host server 102 can also provide services on demand. A module hosting a client is capable of migrating from one server to another server seamlessly, without causing program faults or system breakdown. The host server 102 on the cloud can be managed using a management system.

**[0137]** Other embodiments are within the scope and spirit of the disclosed subject matter.

**[0138]** The subject matter described herein can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structural means disclosed in this specification and structural equivalents thereof, or in combinations of them. The subject matter described herein can be implemented as one or more computer program products, such as one or more computer programs tangibly embodied in an information carrier (e.g., in a machine-readable storage device), or embodied in a propagated signal, for execution by, or to control the operation of, data processing apparatus (e.g., a programmable processor, a computer, or multiple computers). A computer program (also known as a program, software, software application, or code) 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. A computer program does not necessarily correspond to a file. A program can be stored in a portion of a file that holds other programs or data, 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 at one site or distributed across multiple sites and interconnected by a communication network.

**[0139]** The processes and logic flows described in this specification, including the method steps of the subject matter described herein, can be performed by one or more programmable processors executing one or more computer programs to perform functions of the subject matter described herein by operating on input data and generating output. The

processes and logic flows can also be performed by, and apparatus of the subject matter described herein can be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

**[0140]** 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 processor 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 executing 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. Information carriers suitable for embodying computer program instructions and data include all forms of non-volatile memory, 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 optical disks (e.g., CD and DVD disks). The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

**[0141]** To provide for interaction with a user, the subject matter described herein 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.

**[0142]** The techniques described herein can be implemented using one or more modules. As used herein, the term “module” refers to computing software, firmware, hardware, and/or various combinations thereof. At a minimum, however, modules are not to be interpreted as software that is not implemented on hardware, firmware, or recorded on a non-transitory processor readable recordable storage medium. Indeed “module” is to be interpreted to include at least some physical, non-transitory hardware such as a part of a processor or computer. Two different modules can share the same physical hardware (e.g., two different modules can use the same processor and network interface). The modules described herein

can be combined, integrated, separated, and/or duplicated to support various applications. Also, a function described herein as being performed at a particular module can be performed at one or more other modules and/or by one or more other devices instead of or in addition to the function performed at the particular module. Further, the modules can be implemented across multiple devices and/or other components local or remote to one another. Additionally, the modules can be moved from one device and added to another device, and/or can be included in both devices.

**[0143]** The subject matter described herein can be implemented in a computing system that includes a back-end component (e.g., a data server), a middleware component (e.g., an application server), or 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 herein), or any combination of such back-end, middleware, and front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network (“LAN”) and a wide area network (“WAN”), e.g., the Internet.

**[0144]** The terms “a” or “an,” as used herein throughout the present application, can be defined as one or more than one. Also, the use of introductory phrases such as “at least one” and “one or more” should not be construed to imply that the introduction of another element by the indefinite articles “a” or “an” limits the corresponding element to only one such element. The same holds true for the use of definite articles.

**[0145]** It is to be understood that the disclosed subject matter is not limited in its application to the details of construction and to the arrangements of the components set forth in the following description or illustrated in the drawings. The disclosed subject matter is capable of other embodiments and of being practiced and carried out in various ways. Also, it is to be understood that the phraseology and terminology employed herein are for the purpose of description and should not be regarded as limiting.

**[0146]** As such, those skilled in the art will appreciate that the conception, upon which this disclosure is based, may readily be utilized as a basis for the designing of other structures, methods, and systems for carrying out the several purposes of the disclosed subject matter. It is important, therefore, that the claims be regarded as including such equivalent constructions insofar as they do not depart from the spirit and scope of the disclosed subject matter.

**[0147]** Although the disclosed subject matter has been described and illustrated in the foregoing exemplary embodiments, it is understood that the present disclosure has been made only by way of example, and that numerous changes in the details of implementation of the disclosed subject matter may be made without departing from the spirit and scope of the disclosed subject matter.

**[0148]** We claim:

**CLAIMS**

1. An apparatus comprising:
  - a processor configured to run one or more modules stored in memory, wherein the one or more modules are configured to:
    - receive one or more polygons associated with a region of interest;
    - determine a plurality of sub-polygons that are contained within the one or more polygons, wherein each of the sub-polygons is associated with a unique code; and
    - generate a first index system based on at least a subset of the plurality of sub-polygons, thereby providing an efficient mechanism to determine whether a particular location is within the region of interest.
2. The apparatus of claim 1, wherein the unique code includes a location identifier based on a hierarchical encoding scheme.
3. The apparatus of claim 1, wherein the index system includes one of a hash table or a probabilistic data structure.
4. The apparatus of claim 1, wherein the one or more modules are further configured to:
  - determine a difference between the index system and a previously-generated index system; and
  - provide the difference to a computing device to update the previously-generated index system in the computing device.
5. The apparatus of claim 1, wherein the one or more modules are further configured to identify a tree structure that models the unique code of the plurality of sub-polygons based on a hierarchy of the unique codes and generate the first index system based on the tree structure.

6. The apparatus of claim 5, wherein the tree structure includes a branch node and a leaf node, and the branch node is associated with a geographic area that is larger than that of the leaf node, and wherein the one or more modules are configured to generate indices for the index system by traversing the tree structure from the branch node to the leaf node.
7. The apparatus of claim 5, wherein the one or more modules are further configured to:  
  
determine a list of polygons that includes a particular sub-polygon, and  
  
associate the list of polygons with a leaf node of the tree structure corresponding to the particular sub-polygon.
8. The apparatus of claim 5, wherein the one or more modules are further configured to merge the tree structure, corresponding to the region of interest, with a second tree structure corresponding to a second region of interest, thereby providing a single index system modeling both the tree structure and the second tree structure.
9. The apparatus of claim 1, wherein the one or more modules are further configured to remove one or more sub-polygons corresponding to a particular polygon from the index system.
10. The apparatus of claim 1, wherein the one or more modules are further configured to remove one or more sub-polygons from the plurality of sub-polygons to provide a reduced set of sub-polygons and generate the index system from the reduced set of sub-polygons, thereby reducing the size of the index system.
11. The apparatus of claim 1, wherein the one or more modules are further configured to receive a second index system from another computing device and to merge the first index system and the second index system by considering an overlap of sub-polygons corresponding to the first index system and the second index system.

12. The apparatus of claim 1, wherein the one or more modules are further configured to provide the index system to a computing device so that the computing device can use the index system to serve location queries.

13. A method for providing an index system for a region of interest, the method comprising:

receiving, at an index generation module of a computing system, one or more polygons associated with the region of interest;

determining, at the index generation module, a plurality of sub-polygons that are contained within the one or more polygons, wherein each of the sub-polygons is associated with a unique code;

generating, at the index generation module, a first index system based on at least a subset of the plurality of sub-polygons, thereby providing an efficient mechanism to determine whether a particular location is within the region of interest.

14. The method of claim 13, wherein the index system includes one of a hash table or a probabilistic data structure.

15. The method of claim 13, further comprising:

determining, by the index generation module, a difference between the index system and a previously-generated index system; and

providing, by the index generation module, the difference to a computing device to update the previously-generated index system in the computing device.

16. The method of claim 13, further comprising identifying a tree structure that models the unique code of the plurality of sub-polygons based on a hierarchy of the unique codes and generate the first index system based on the tree structure.

17. The method of claim 16, wherein the tree structure includes a branch node and a leaf node, and the branch node is associated with a geographic area that is larger than that of the leaf node, and wherein the method further comprises generating, at the index generation

module, indices for the index system by traversing the tree structure from the branch node to the leaf node.

18. The method of claim 16, further comprising merging the tree structure, corresponding to the region of interest, with a second tree structure corresponding to a second region of interest, thereby providing a single index system modeling both the tree structure and the second tree structure.

19. The method of claim 13, further comprising receiving a second index system from another computing device and merging the first index system and the second index system by considering an overlap of sub-polygons corresponding to the first index system and the second index system.

20. The method of claim 13, further comprising providing the index system to a computing device so that the computing device can use the index system to serve location queries.

21. A non-transitory computer readable medium having executable instructions operable to cause a data processing apparatus to:

- receive one or more polygons associated with a region of interest;
- determine a plurality of sub-polygons that are contained within the one or more polygons, wherein each of the sub-polygons is associated with a unique code; and
- generate a first index system based on the tree structure, thereby providing an efficient mechanism to search whether a particular location is within the region of interest.

22. The computer readable medium of claim 21, wherein the index system includes one of a hash table or a probabilistic data structure.

23. The computer readable medium of claim 21, wherein the executable instructions are further operable to cause the data processing apparatus to:

determine a difference between the index system and a previously-generated index system; and

provide the difference to a computing device to update the previously-generated index system in the computing device.

24. The computer readable medium of claim 21, wherein the executable instructions are further operable to cause the data processing apparatus to identify a tree structure that models the unique code of the plurality of sub-polygons based on a hierarchy of the unique codes and generate the first index system based on the tree structure.

25. The computer readable medium of claim 24, wherein the tree structure includes a branch node and a leaf node, and the branch node is associated with a geographic area that is larger than that of the leaf node, and wherein the method further comprises generating, at the index generation module, indices for the index system by traversing the tree structure from the branch node to the leaf node.

26. The computer readable medium of claim 21, wherein the executable instructions are further operable to cause the data processing apparatus to merge the tree structure, corresponding to the region of interest, with a second tree structure corresponding to a second region of interest, thereby providing a single index system modeling both the tree structure and the second tree structure.

27. The computer readable medium of claim 21, wherein the executable instructions are further operable to cause the data processing apparatus to receive a second index system from another computing device and merge the first index system and the second index system by considering an overlap of sub-polygons corresponding to the first index system and the second index system.

28. An apparatus comprising:  
one or more interfaces configured to provide communication with a client device via a communication network;

a memory device that is configured to maintain a hierarchical index system of unique identifiers corresponding to sub-polygons in a region of interest; and

a processor, in communication with the one or more interfaces, and configured to run one or more modules that are configured to:

receive a location query from the client device, wherein the location query includes a location identifier associated with the client device;

determine a query identifier corresponding to the location identifier;

compare the query identifier with the index system to determine that the location identifier provided by the client device is within the region of interest; and

provide a service associated with the region of interest to the client device over the communication network.

29. The apparatus of claim 28, wherein the unique identifiers and the query identifier include location identifiers based on a hierarchical encoding scheme.

30. The apparatus of claim 29, wherein the one or more modules are configured to compare one or more bits in the query identifier with bits in the index system to determine that the query identifier is represented in the index system and that the location identifier provided by the client device is within the region of interest

31. The apparatus of claim 30, wherein the one or more modules are configured to compare a first set of bits of the query location identifier, corresponding to a lower-precision sub-polygon, to the index system before comparing a second set of bits of the query location identifier, corresponding to a higher-precision sub-polygon.

32. The apparatus of claim 31, wherein the index system comprises an index tree, when the first set of bits of the query identifier matches a first index of the index system corresponding to a leaf node of the index tree, the one or more modules that are configured to determine that the query identifier is within the region of interest.

33. The apparatus of claim 28, wherein the one or more modules are configured to retrieve, from the index system, a polygon identifier associated with the query identifier, determine a group identifier associated with the polygon identifier, and provide the service associated with the group identifier to the client device over the communication network.

34. The apparatus of claim 28, wherein the one or more modules are configured to retrieve, from the index system, a polygon identifier associated with the query identifier, and provide data associated with the polygon identifier to the client device over the communication network.




FIG. 1

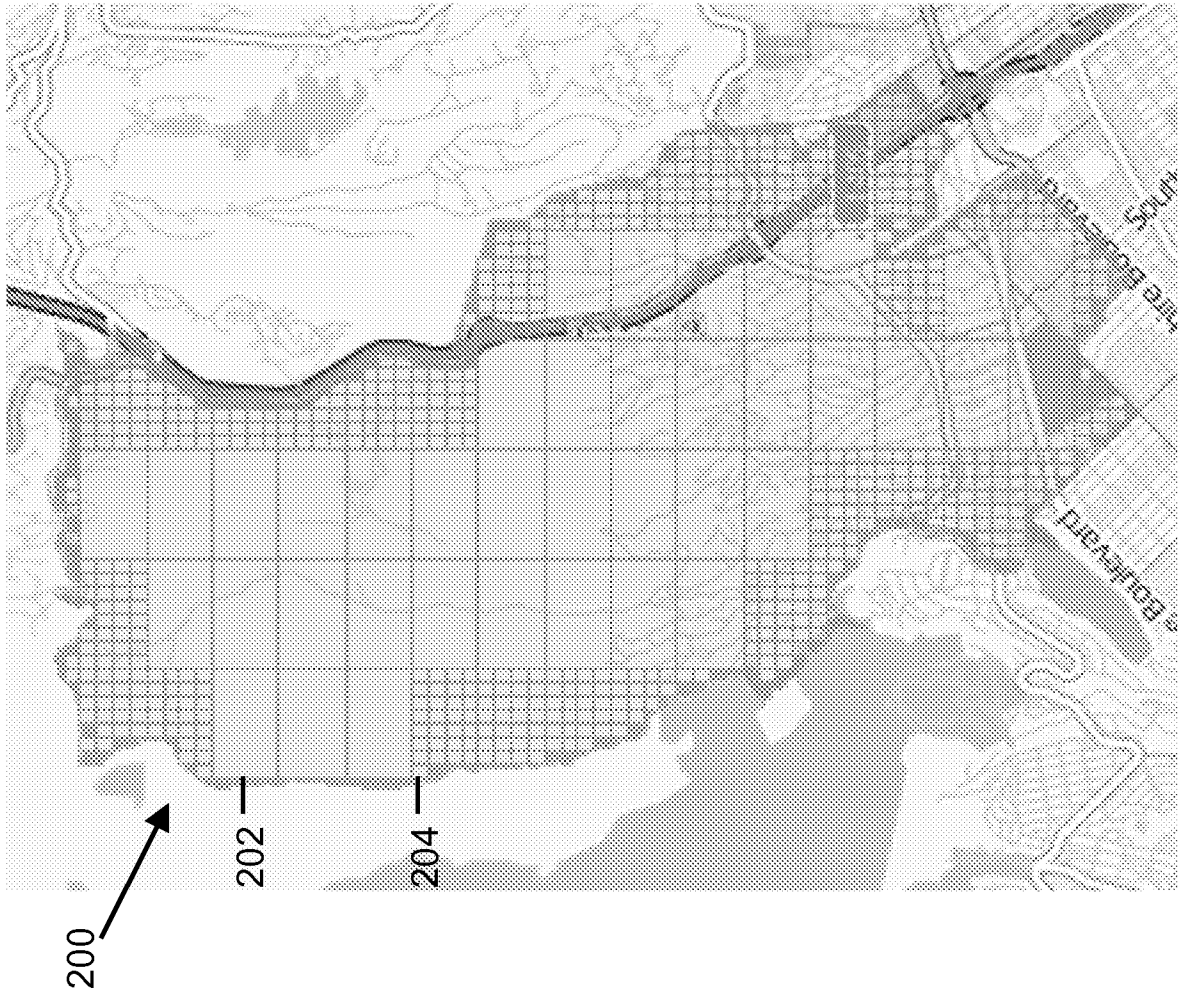


FIG. 2






FIG. 4




FIG. 5






FIG. 6



(A)



(B)



(C)

FIG. 7




FIG. 8A



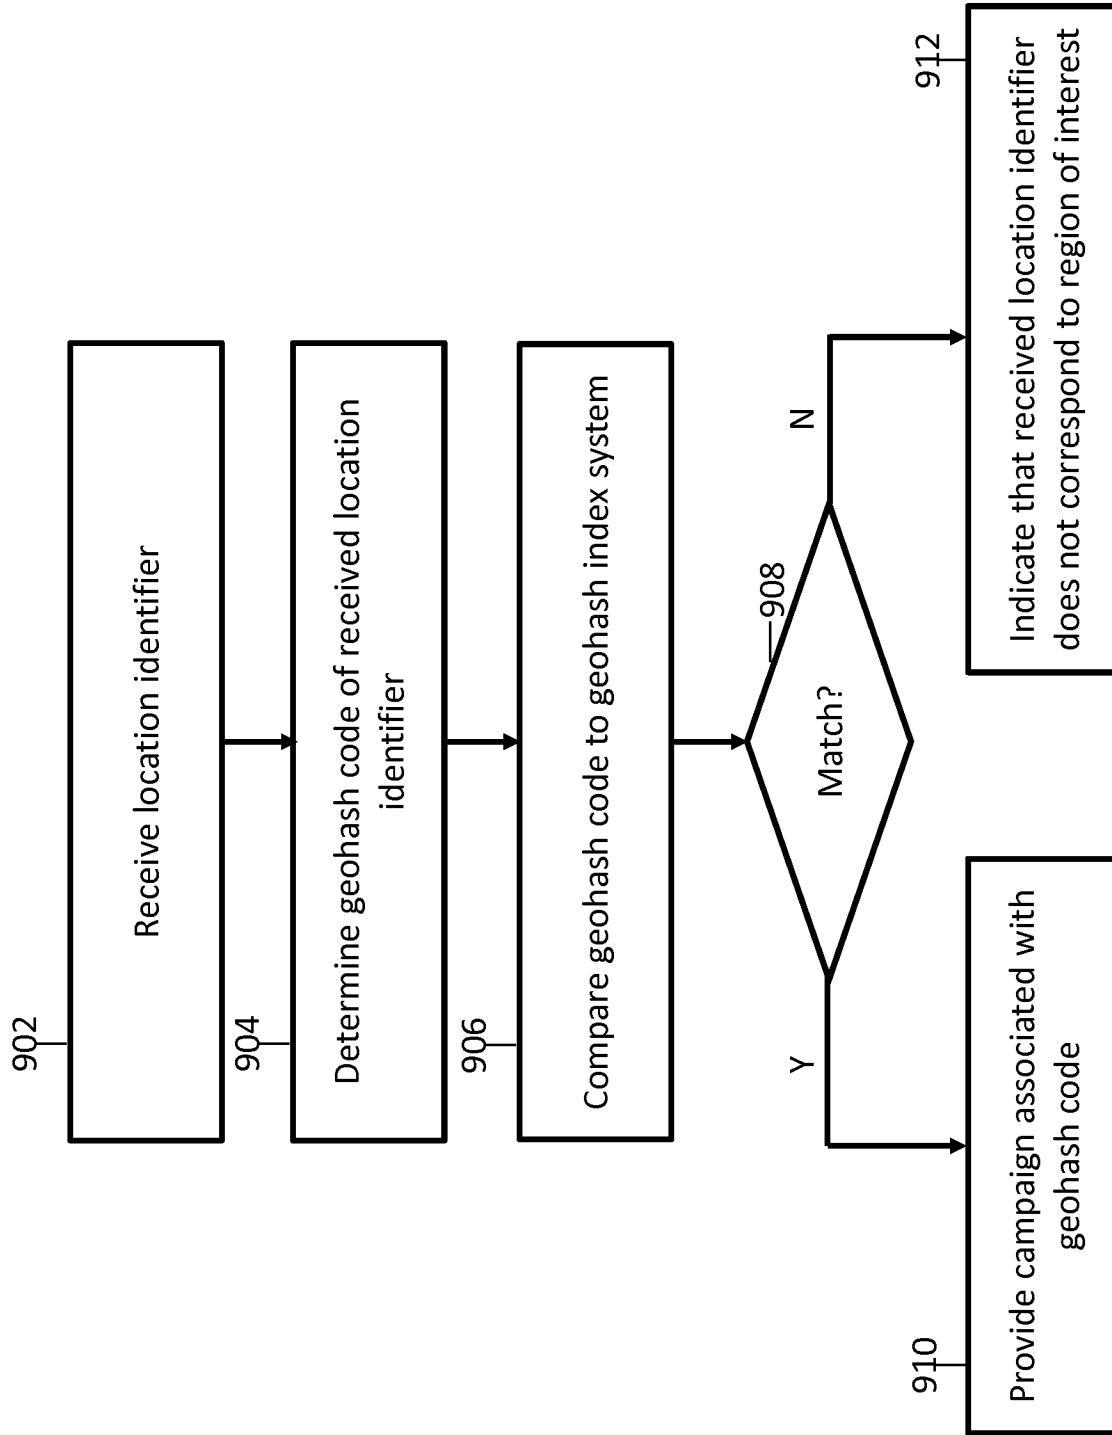


FIG. 9




FIG. 10




FIG. 11

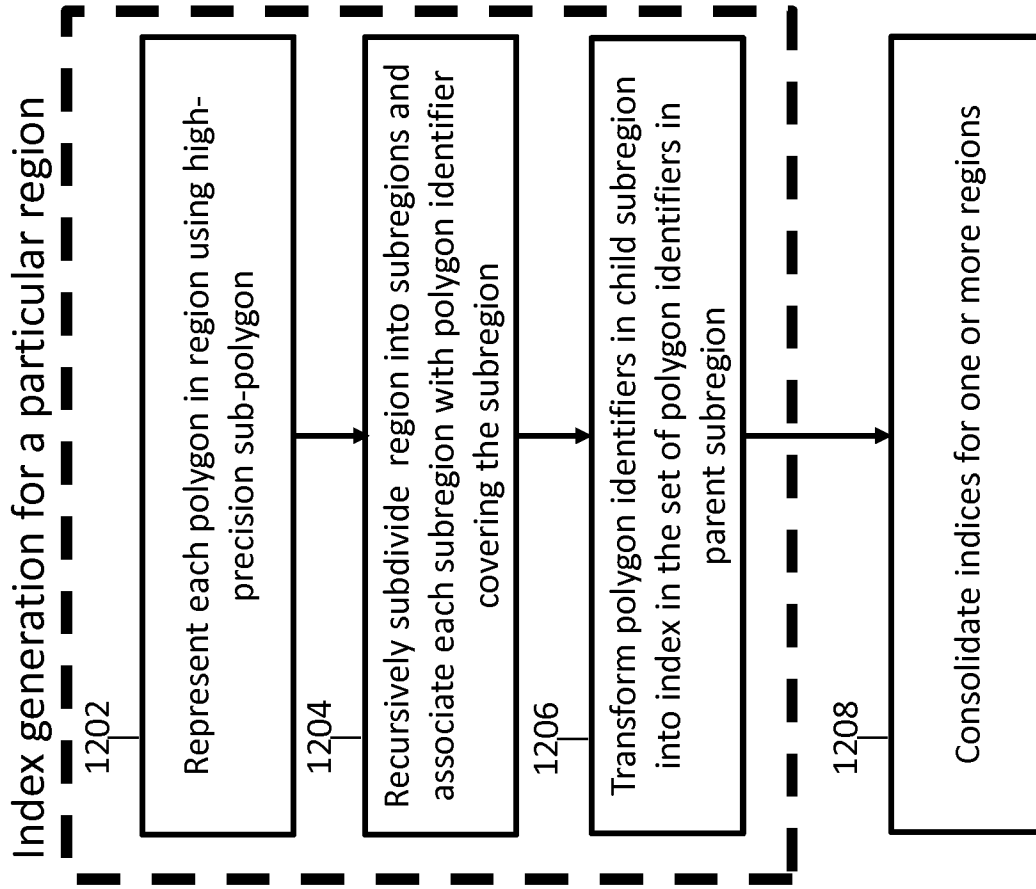


FIG. 12






FIG. 13B




FIG. 13C

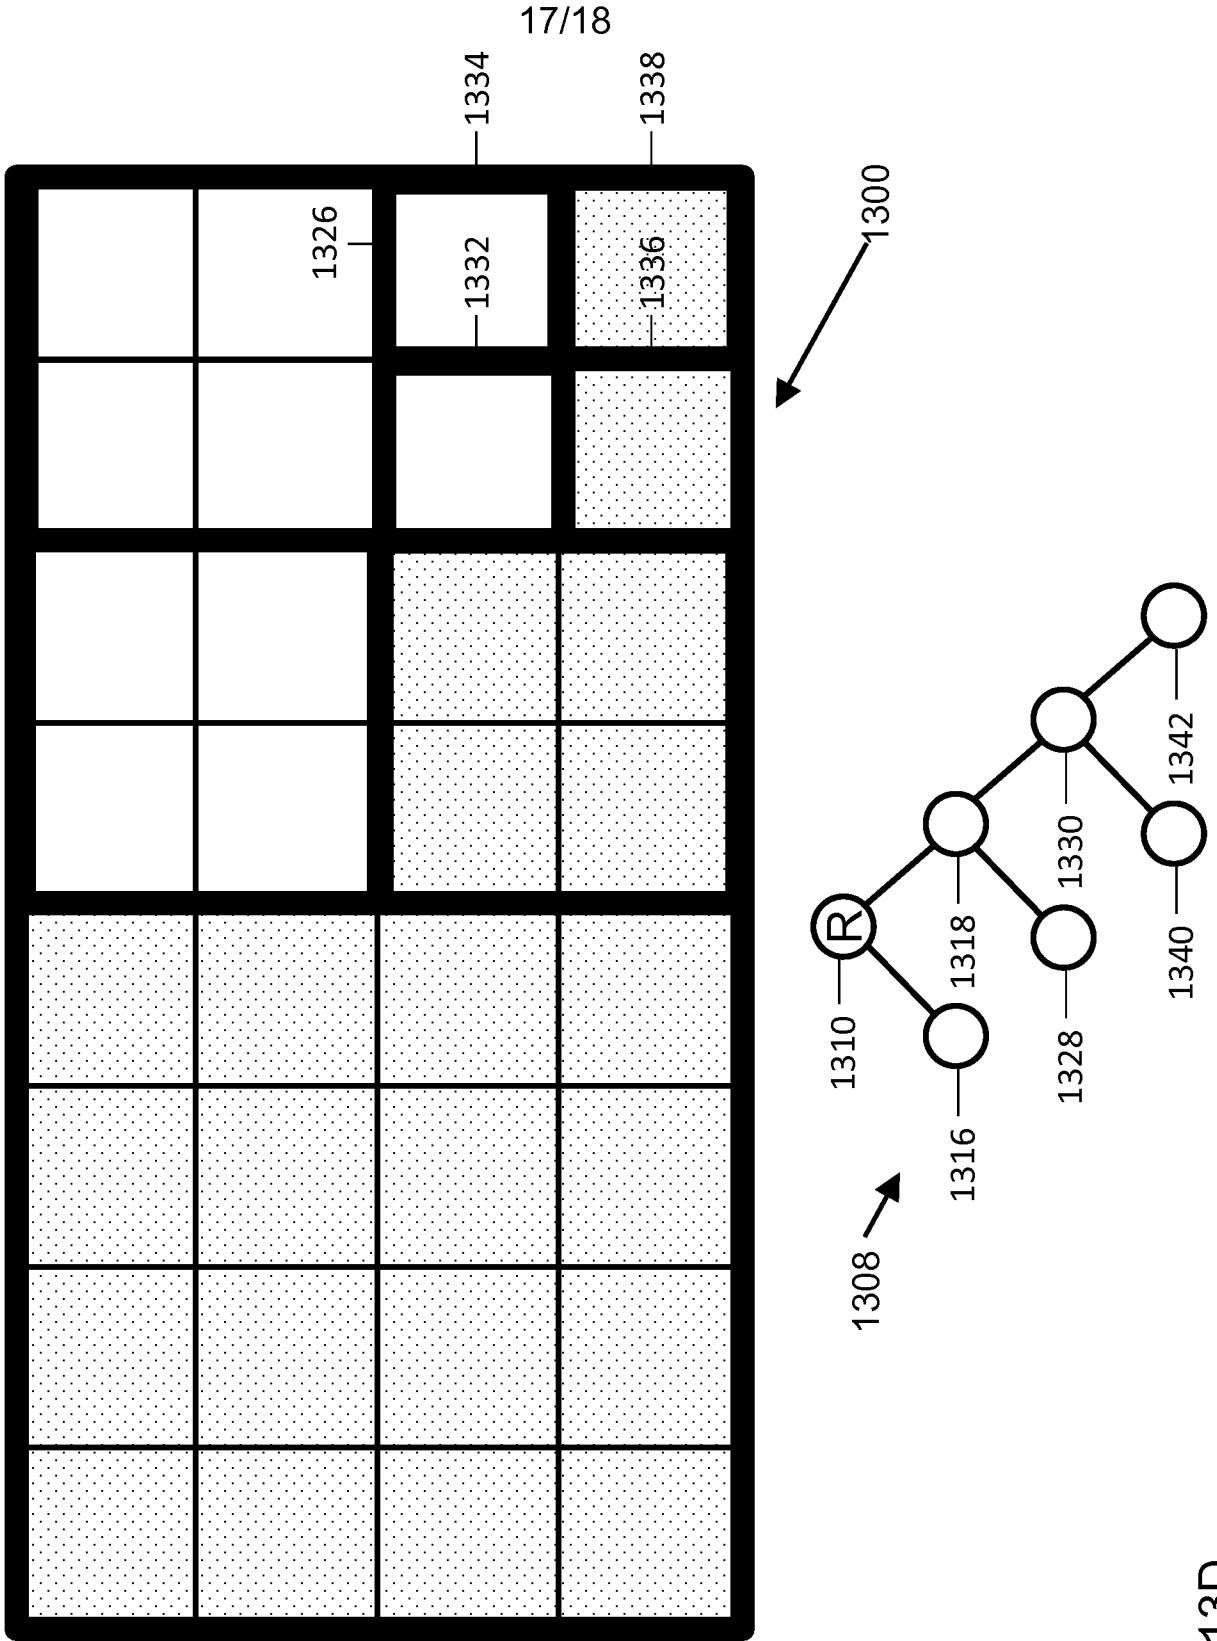


FIG. 13D

1400 ↗




FIG. 14

INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2014/029724

A. CLASSIFICATION OF SUBJECT MATTER  
INV. G06F17/30  
ADD.  
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED  
Minimum documentation searched (classification system followed by classification symbols)  
G06F

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Anonymous: "SpatialPrefixTree (Lucene 4.0.0 API)", 13 March 2013 (2013-03-13), XP055129916, Retrieved from the Internet: URL:https://web.archive.org/web/20130313072002/http://lucene.apache.org/core/4_0_0/spatial/org/apache/lucene/spatial/prefix/tree/SpatialPrefixTree.html [retrieved on 2014-07-17] the whole document ----- -/--	1-34

Further documents are listed in the continuation of Box C.

See patent family annex.

\* Special categories of cited documents :

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

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

Date of the actual completion of the international search  18 July 2014	Date of mailing of the international search report  29/07/2014
---	--

Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer  Denoual, Matthieu
--	---

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2014/029724

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	David Smiley: "Lucene 4 spatial", Open Source Search Conference, 2 October 2012 (2012-10-02), pages 1-16, XP055129458, Chantilly, Virginia, USA Retrieved from the Internet: URL:http://info.basistech.com/Portals/1654 17/presentations/oss-2012/lucene-4s-new-spatial-module-david-smiley-open-source-search-conference-2012.pdf?&_hssc=&_hstc&hsctaTracking=068335f3-cf17-4140-8012-b59a2cef5ab9 ba7659a3-2c38-4e71-9ef8-61d952ae58d1 [retrieved on 2014-07-16] the whole document	1-34
X	----- SAMET H ED - ALEJANDRO P BUCHMANN ET AL: "Hierarchical Spatial data structures", DESIGN AND IMPLEMENTATION OF LARGE SPATIAL DATABASES. FIRST SYMPOSIUM SSD '89, SPRINGER, DE, 17 July 1989 (1989-07-17), pages 193-212, XP002566914, ISBN: 978-3-540-52208-9 the whole document	1-34
X	----- US 2006/253481 A1 (GUIDO MATTHEW A [US] ET AL) 9 November 2006 (2006-11-09) the whole document	1-34
X	----- OOSTEROM VAN P ET AL: "The Spatial Location Code", PROCEEDINGS OF THE INTERNATIONAL SYMPOSIUM ON SPATIAL DATAHANDLING, XX, XX, 12 August 1996 (1996-08-12), pages 3B.1-3B.17, XP002293525, the whole document	1-34
X	----- VARMA H P ET AL: "A DATA STRUCTURE FOR SPATIO-TEMPORAL DATABASES", INTERNATIONALE HYDROGRAPHIC REVIEW, INTERNATIONAL HYDROGRAPHIC BUREAU, MONACO, MC, vol. 67, no. 1, 1 January 1990 (1990-01-01), pages 71-92, XP000565997, ISSN: 0020-6946 the whole document	1-34
A	----- US 7 576 754 B1 (JOSEPH JOSHY [US] ET AL) 18 August 2009 (2009-08-18) the whole document	1-34
	-----	

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2014/029724

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006253481	A1	09-11-2006	
		EP 1271458 A2	02-01-2003
		ES 2459201 T3	08-05-2014
		JP 4490028 B2	23-06-2010
		JP 2003108591 A	11-04-2003
		US 7089264 B1	08-08-2006
		US 2006253481 A1	09-11-2006
-----			
US 7576754	B1	18-08-2009	NONE
-----			