(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
3 January 2014 (03.01.2014)

WIPO | PCT

(10) International Publication Number
WO 2014/004741 A1

(51) International Patent Classification:
G06Q 30/02 (2012.01)

(21) International Application Number:
PCT/US2013/048006

(22) International Filing Date:
26 June 2013 (26.06.2013)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/664,734 26 June 2012 (26.06.2012) US

(71) Applicant: **MEDIO SYSTEMS, INC.** [US/US]; One Convention Place, 701 Pike Street, 15th Floor, Seattle, Washington 98101 (US).

(72) Inventors: **LANPHEAR, Damon**; One Convention Place, 701 Pike Street, 15th Floor, Seattle, Washington 98101 (US). **BISWAS, Prabuddha**; One Convention Place, 701 Pike Street, 15th Floor, Seattle, Washington 98101 (US).

(74) Agent: **HALL, David**; Kilpatrick Townsend & Stockton LLP, Two Embarcadero Center, Eighth Floor, San Francisco, California 94111 (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, 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, 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))*

(54) Title: DISTRIBUTED, REAL-TIME ONLINE ANALYTICAL PROCESSING (OLAP)

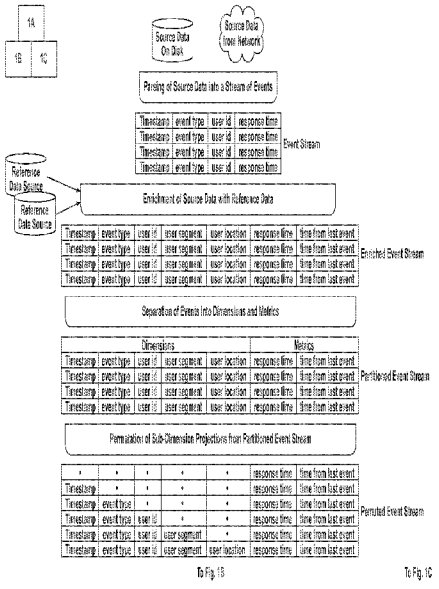


FIG. 1A

(57) Abstract: Source data of an event stream is parsed and supplemented with additional data from reference data sources, producing an enriched event stream from the parsed event stream data. The data records of the enriched event stream are partitioned into data fields designated as a dimension partition and a metric partition, which are partitioned into sub-dimension projections mapped to a plurality of storage keys, such that each of the storage keys includes one or more placeholder wildcard values and each of the storage keys is stored into a database of the computer system by the computer processor. The stored storage keys are then aggregated onto a two-dimensional coordinate vector such that, if the computer processor identifies a permuted storage key having metric values for which a corresponding storage key already exists in the database, then the computer processor aggregates the metric values of the identified storage key with the metric values of the corresponding storage key, and if the computer processor does not identify the permuted storage key as having a corresponding storage key that already exists in the database, then the computer processor writes the metric values of the permuted storage key into the database, comprising initial values for the key combination of dimension values.

WO 2014/004741 A1

DISTRIBUTED, REAL-TIME Online Analytical Processing (OLAP)

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority of co-pending U.S. Provisional Patent Application Serial No. 61/664,734 entitled "DISTRIBUTED, REAL-TIME Online Analytical Processing (OLAP)" by Damon Lanphear and Prabuddha Biswas, filed June 26, 2012. Priority of the filing date of June 26, 2012 is hereby claimed, and the disclosure of the Provisional Patent Application is hereby incorporated by reference

BACKGROUND

[0002] Online Analytical Processing (OLAP) is a set of techniques for organizing source data to support rapid responses to complex, multi-dimensional queries of the source data. OLAP techniques are commonly applied to business intelligence and marketing problems in which it is desired to investigate how customers behave according to demographic, psychographic, or other attributes. Other applications of OLAP involve relational reporting and data mining. The output of an OLAP query is typically provided as a table having rows and columns comprising "dimensions" and values of the table cells comprising "measures". The dimensions may be categorized by searchable terms such as sales, date, clerk, customer name, and the like. The measures comprise data values that correspond to the intersection of particular combinations of the dimensions. For example, a query might request date = May 1st, 2011 and clerk = John Smith. The measures would correspond to all data records where the sales date is May 1st and the clerk name is John Smith.

[0003] The set of all possible combinations of the source data dimensions contain the set of responses to all possible queries submitted to the OLAP system. To ensure timely response to OLAP queries, and because of the large number of possible queries that could be submitted, many OLAP systems calculate a predetermined number of the possible OLAP queries. If a query is submitted that matches a pre-calculated response, then the response table can be provided immediately. Typically, queries that are most likely to be received are the subject of pre-calculation. Responses to other queries must be solved on demand, upon receipt.

[0004] Efficient processing for responding to queries involves selecting which dimension combinations, or views, are most likely to be requested from users in a query, and pre-calculating responses for rapid query response. Other techniques for OLAP involve combining use of relational tables and specialized storage for improved response. It is known to provide what is called "real-time OLAP", in which the OLAP database (commonly called the "cube") is stored in RAM, so that responses to OLAP queries can be calculated on the data stored in RAM and returned in real-time. In this type of OLAP, the size of the OLAP data for searching can be no larger than the size of RAM on the computer machine that hosts the OLAP application.

[0005] Several approaches to implementing OLAP exist today and have been in the commercial marketplace for over a decade. As apparent from the discussion above, the existing approaches generally suffer from fundamental limitations on the amount of data that can be searched, the number of dimensions in the database, cardinality of the dimensions, and ability to be updated in real-time as new data become available. These limitations render the existing approaches impractical for new classes of OLAP applications that include live dashboards on customer activity and predictive analytical features.

SUMMARY

[0006] In accordance with the present invention, source data of an event stream is parsed and supplemented with additional data from reference data sources, producing an enriched event stream from the parsed event stream data. The data records of the enriched event stream are partitioned into data fields designated as a dimension partition and a metric partition. The dimension partition and the metric partition are partitioned into sub-dimension projections that are mapped to a plurality of storage keys, such that each of the storage keys includes one or more placeholder wildcard values and each of the storage keys is stored into a database of the computer system by the computer processor. The stored storage keys are then aggregated onto a two-dimensional coordinate vector such that, if the computer processor identifies a permuted storage key having metric values for which a corresponding storage key already exists in the database, then the computer processor aggregates the metric values of the identified storage key with the metric values of the corresponding storage key, and if the computer processor does not identify the permuted storage key as having a corresponding storage key that already exists in the database, then the computer processor writes the metric values of the permuted storage key into the database, comprising initial values for the key combination of dimension values. The resulting database increases the efficiency of operations such that the system experiences an

increase in the amount of data that can be quickly searched, the number of data dimensions that can be processed in the database, the cardinality of the dimensions, and the ability to be updated in real-time as new data become available.

[0007] In another aspect, as disclosed herein, a distributed, persistent data store is provided that is structured as an associative array. The associative array may span a number of computer hosts, and new computer hosts can be added to the array at any time to expand its capacity. The distributed associative array may include a two-level caching mechanism by which data are stored first on a persistent medium of a host computer and secondarily cached in volatile memory of the host computer for fast access. The distributed associative array supports atomic updates to its contents, providing guarantees on the coherency of the data in the presence of concurrent updates. This construction provides a system with theoretically unlimited scalability in the size of the source data, number of dimensions, and the cardinality of the dimensions. Further, queries and updates to the OLAP system can be processed in an amount of time that is sufficiently rapid to be useful for interactive applications.

[0008] Other features and advantages of the present invention will be apparent from the following description of the embodiments, which illustrate, by way of example, the principles of the invention.

[0009] Additional details are provided by the attached appendices, which are incorporated herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIG. 1 is a process flow diagram that illustrates data operations in a system constructed in accordance with the invention.

[0011] FIG. 2 is a block diagram of the system topology for the system illustrated in FIG. 1.

[0012] FIG. 3 is a workflow diagram for the system illustrated in FIG. 1.

[0013] Figure 4 is a block diagram of an end user mobile device such as used by an end user of the system.

[0014] Figure 5 is a schematic diagram representation of a computer system for implementing functions and operations as described herein.

DETAILED DESCRIPTION

[0015] FIG. 1 illustrates the process flow that takes place in a system constructed in accordance with the invention. FIG. 1 illustrates aspects of the system, including a write path, a read path, and a request routing operation for request routing between the read and write paths. More particularly, the write path of FIG. 1 shows that source data is available from storage that is local to an OLAP host system, or over a network, in a write path that populates an OLAP database of the system (the write path that begins in the upper left corner of FIG. 1). In the read path, starting in the upper right corner of FIG. 1, a data query message is received at the system, as indicated by the network cloud at the top right corner of FIG. 1. The query is received and processed by the OLAP Query Engine of the system, and then is processed into the proper format and then the query is applied from the read path to the system database at the bottom of FIG. 1. The results of the query are returned to a results merge operation in accordance with request routing, back up through the read path for return to the query initiator at the upper right of the read path in FIG. 1.

[0016] The write path of FIG. 1 shows that the source data is parsed into a stream of events, such as sales events or service calls, using appropriate dimensions of interest to system users. For example, in FIG. 1, each source event is indicated as a 4-tuple, comprising a data record that includes a time stamp value, an event type indicator, a user ID, and a response time. Such data records might be produced by a system that keeps track of service calls or product complaints. In such a system, the "event type indicator" value could be used to distinguish between service calls, repair requests, product complaints, and the like. In a point-of-sales system, the source data records would more likely data relating to sales by persons for which the data includes items and pricing. Parsing such data into an event stream would produce tuples that are combinations of, for example, a time stamp, sales person ID, price, and number of items sold. For purposes of illustration, a sample 4-tuple value from an event stream comprising point-of-sales events might be given by the following:

Event Stream:

[TimeStamp, SalesPersonID, Price, Items]

[2012-01-19T13:34:00-0800, 11923245, $2352, 40].

[0017] In the sample Event Stream 4-tuple data record listed immediately above, the time stamp value may be read as year=2012, month=01, day=19, and time=13:34:00-0800, wherein

time is given in a 24-hour format relative to Greenwich mean time. Other formats may be used without departing from the teachings of the invention.

[0018] The next operation produces an enriched event stream from the parsed event stream data. In this operation, reference data sources may be consulted for adding additional data of interest to each parsed event stream data, according to additional parameters that may be of interest and that may be accessed via the OLAP query system. For example, a sales person may have a particular store location that may be determined from a reference data source for adding to the event stream data. Such additional data is illustrated in FIG. 1 as comprising data relating to user segment data, user location data, and time-from-last-event data. Thus, the 4-tuple of the parsed event stream data is converted to the 7-tuple illustrated in FIG. 1. In another embodiment, continuing with the example of an event stream comprising point-of-sales events, an exemplary enriched event stream data record might be given by the following:

Enriched Event Stream:

[TimeStamp, SalesPerson, Location, Price, Items]

[2012-01-19T13:34:00-0800, John, Seattle, $2352, 40]

[0019] In the sample Enriched Event Stream 7-tuple data record listed immediately above, the event stream data record is enriched by substitution of the sales person ID ("11923245") by a SalesPerson name data value ("John") and by the addition of a sales location ("Seattle"). Thus, the event stream data is enriched with reference data that provides additional data of importance to the system user. The reference data may comprise, for example, data about the entity that created the corresponding source data, such as the user geographic location or demographic (segment) information regarding the customer whose action generated the source data.

[0020] Next in the write path, the incoming data is partitioned into events according to dimensions and measures, referred to as metrics. The metrics generally relate to factual data points such as time elapsed, monetary amounts, street addresses, client age, and the like. The dimensions generally relate to attributes that identify the incoming data records, such as timestamp for time of creation, user demographic group, transaction type, and the like. FIG. 1 shows that the data records of the enriched event stream are partitioned into data fields designated as dimensions and data fields designated as metrics. In FIG. 1, the dimension fields include timestamp, event type, user ID, user segment, and user location, whereas the metric fields include response time and time from last event. In another embodiment, continuing with

the example of an event stream comprising point-of-sales events, an exemplary partitioned event stream data record might be given by the following:

Partitioned Event Stream:

[Dimensions: Year, Month, Day, SalesPerson, Location | Metrics: Price, Items]

[2012, 01, 19, John, Seattle, 2352, 40]

[0021] In the sample Partitioned Event Stream data record immediately above, the data field schema is given by the 5-tuple comprising dimensions of <Year>-<Month>-<Day>-<SalesPerson>-<Location>, and the metrics of interest comprise the 2-tuple of metrics <sum(Price)>-<sum(Items)>. Thus, the timestamp value is converted to more meaningful year-month-day values. Other data conversions and transformations may be utilized, as desired for the OLAP query system.

[0022] The partitioned event stream is then permuted into sub-dimension projections, as will be known to those skilled in the art. The permuted sub-dimension projections are then mapped to storage keys, again, as will be known to those skilled in the art. For example, the permutation of dimensions may be performed by a processing function that generates a permutation of dimensions and a wildcard placeholder to represent an "ANY" value for a dimension. The ANY value is an aggregation of all metrics that match a particular dimension. The system can utilize the permutations, for example, to enable a user to ask for "total value for all sales in Seattle for all time," or to ask for "all sales for the month of July made by salesperson John." As known to those skilled in the art, such permutation processing will permute all combinations of possible ordered subsets that will allow lookups by any key combinations of the dimensions and metrics.

[0023] With respect to the permutation processing, and continuing with the example of an event stream comprising point-of-sales events, an exemplary list of permutations of the 7-tuple partitioned event stream records with dimensions and metrics might be given by the following:

Permuted Keyed Event Stream:

[Year.Month.Day.Location.SalesPerson, Price, Quantity]

[*.*.*.*.* , 2352 , 40],

[2012.*.*. *.* , 2352, 40],

[2012.01.*.*.*, 2352, 40],

[2012.01.19.*.*, 2352, 40],

[2012.01.19.Seattle.*, 2352, 40],

[2012.01.19.Seattle.John, 2352, 40],

[*.01.19.Seattle.John, 2352, 40],

[2012.*.19.Seattle.John, 2352, 40],

[2012.01.*.Seattle.John, 2352, 40],

[2012.01.19.*.John, 2352, 40],

... and so forth for the remaining permutations, as will be known to those skilled in the art. Each of these permuted data records comprises a key for a mapping of the partitioned event stream data to the collection of permuted event stream data. As each permuted data record is generated, it is stored in a system database. Each permuted data record comprises a data cell of an OLAP cube.

[0024] As the permuted data records are generated and stored, the system processing in the write path uses aggregation rules to apply aggregate functions to the incoming data. The aggregate functions map the cube data cells onto a two-dimensional coordinate vector. During aggregation, if a key (i.e., permuted data record) already exists in the database, then the metric values of the data record are aggregated with the existing values in the database according to the aggregation function. If the key does not already exist, then the metrics are written into the database as the initial values for that key combination of dimension values. This technique for storage and processing of event stream data allows efficient OLAP analysis on the data fields as desired. For example, continuing with the example of an event stream comprising point-of-sales events, the OLAP analysis system as disclosed herein supports query processing that can retrieve data records that satisfy a query asking for total items sold and total revenue received by all combinations of year, month, day, location, and sales person.

[0025] The storing of the permuted data records and the aggregating according to identified already existing metric values for mapping onto a two-dimensional coordinate vector provides a distributed, persistent data store that may be structured as an associative array. Those skilled in the art will understand that conventional OLAP cubes typically have three or more data dimensions, over which database processing and query processing must be provided. The mapping of n-tuple permuted data records onto a two-dimensional coordinate vector as disclosed

herein increases efficiency by decreasing the dimensions of the OLAP cube over which processing must be provided.

[0026] The associative array may be stored in the database of the system and can span across a number of computer hosts. Moreover, new computer hosts can be added to the associative array at any time to expand its capacity. In this way, the associative array can be distributed across multiple computer hosts. Thus, the distributed associative array can be provided with a two-level caching mechanism, by which data are stored first on the persistent medium of the host computer and secondarily cached in volatile memory for fast access. The distributed associative array supports atomic updates to its contents, providing guarantees on the coherency of the data in the presence of concurrent updates.

[0027] The query technique that is used for the system can comprise a wide variety of techniques. For example, there may be many query front ends that are designed to interface with difference systems. Such systems are suitable for use with the disclosed system. In general, the fundamental concept of satisfying a query requires that the query language must be translated to the internal vector representation used by the analysis system. In some cases, the keys may be accessed directly, and in other cases an intermediary language such as "MDX" may be utilized to access the keys (permuted data records).

[0028] The "Computer 1" and "Computer N" blocks at the bottom of FIG. 1 indicate the computers that store the OLAP cube data. That is, FIG. 1 indicates that one or more computers will receive incoming data processed according to the write path, and can be used to respond to OLAP queries from one or more clients in the read path. The collection of persistent data storage at each of the computer blocks in FIG. 1 represents a collection of computer hosts that provide an associative array that are distributed across the computer blocks Computer 1, ..., Computer N. The collection of computer hosts may be managed, for example, by one of the host computers that is designated a managing host computer. Alternatively, the collection of computer hosts may distribute incoming event stream data for processing according to predetermined techniques for appropriate load balancing. New computer hosts can be added to the associative array at any time to expand its capacity.

[0029] The distributed associative array may include a two-level caching mechanism by which data are stored first on the persistent medium of one or more of the host computers and secondarily cached in volatile memory of the host computers for fast access. The two-level caching is achieved by first storing the data in a distributed file system. The distributed file

system keeps three copies of the data for redundancy. The system then assigns a portion of the total key-space to each host computer, which acts as a server of the database, such that each search has 1/N of the total key space, where N is the number of host computer servers. The assignment of key-space across the host computers of the system may be achieved under management of a designated managing host computer or the host computers may operate according to a predetermined collaborative assignment scheme. The key assignment is based on a lexicographic sorting of the keys and 1/N of that key space is assigned at random to each of the servers. The servers will then cache their portion of the key space in memory. The distributed associative array supports atomic updates to its contents, providing guarantees on the coherency of the data in the presence of concurrent updates.

[0030] Each of the host computers Computer 1, ..., Computer N includes a network communications block adapted for communications with a computer network, such as the Internet or a local area network of host computers, and also includes a computer processor and associated memory for use by the computer processor. Those skilled in the art will understand additional physical components that may be necessary for, or may improve the processing of, the host computers.

[0031] In the read path, starting in the upper right corner of FIG. 1, a data query message is received in accordance with an OLAP query client protocol, as indicated by the network cloud at the top right corner of FIG. 1. The query is received and processed by the OLAP Query Engine of the system, and then is processed (or translated) to include the appropriate storage key information according to the sub-dimension mapping referenced above in the write path. The query is then applied to the database, such that the query may be provided to the distributed computers of the OLAP system, comprising system computers Computer 1, Computer 2, and so forth to Computer N, as illustrated in FIG. 1. The results of the query are returned to a results merge operation in accordance with request routing, so that results from each of the system computers is merged in the Query Execution and Result Merge block, to be returned to the querying client at the upper right of the read path in FIG. 1.

[0032] The operation of the write path and read path of the OLAP system described herein will be better understood with accordance with the following description of the OLAP system.

[0033] OLAP data is commonly referred to as a cube, where each dimension of the cube corresponds to an attribute of the source data. For example, a customer's location at the time of purchase can be regarded as a dimension. Another dimension may correspond to the purchased

item, another dimension may correspond to the user's demographic attributes, and so forth. The OLAP data is represented in the associative array as a two-dimensional projection of a N-dimensional cube. The cube is projected into two-dimensions by composing a coordinate vector of a data point in the cube and applying this vector as a key in the associative array. Each array key is associated with a value, which is a value of the data point. The data point value corresponds to one or more metrics for the given cube coordinate. For example, the data point value may be the number of purchases, and the dimensions may comprise the geographic location, demographic segment, and time of day. Thus, if the coordinate vector comprising the two-dimensional projection is represented as $\{a, b\}$, then the two dimensions of the example coordinate vector correspond to a = [number of purchases] and b = [geographic location, demographic segment, time of day] where b is a 3-tuple.

[0034] Another example of the data organization comprising an associative array represented as a two-dimensional projection of a N-dimensional cube is illustrated as follows. Taking sales data for ringtones as the example, suppose that the data is formatted as a 5-tuple comprising {timestamp, userid, geographic location, ringtoneid, sale-amount}. An example of a data record with actual values for the five items of the 5-tuple might be: {2011-03-30T23:13:23, X1EhipHMjNP5ZQ, Seattle, 4a62b1170c77, $0.99}.

[0035] If we assume, for the example, that the OLAP cube according to the invention includes data (i.e., metrics) for Total Sales Amount and Number of Sales, then the system receives data records corresponding to the 5-tuple above and aggregates the data according to the database aggregation rules. That is, metrics for the OLAP cube will include Total Sales Amount and Number of Sales. Since the data records include date information and geographic location information, the dimensions on which we may query include Month and Location. The system described herein will permute each dimension so all combinations of dimensions are included and can be specified by a wildcard value, represented by "*", meaning that any value for the dimension may be searched. Each dimension is represented as a tuple of the dimension values, for example, the permutations of dimension values month and location are represented in the table below:

.
March.*

*.Seattle
March.Seattle

The aggregated metrics then provide one or more columns for each dimension, corresponding to the aggregated metrics. For example, for the OLAP metrics of total Sales Amount and Number of Sales, the following table is provided in the database, where for each dimension representation (comprising four rows for aggregate of all months and locations, month = March, location = Seattle, and both month = March and location = Seattle) there are one or more columns corresponding to the aggregated metrics, with Total Sales in the second column, and Number of Sales in the third column:

.	$126863.55	128145
March.*	$1450.00	1525
*.Seattle	$65159.00	67145
March.Seattle	$531	615

[0036] In the example above, the data includes two dimensions, comprising rows and columns. The rows associate the aggregate values with a permutation of the key space, and the columns represent the values of aggregates for the given key. Thus, the key space may be regarded as an N-dimensional matrix where each dimension corresponds to a dimension of the data. In this example, from the table above, those dimensions would be Time and Place. For the purpose of discussion, imagine that time is the x-axis, and place is the y-axis. The aggregate value would be determined by first finding time on the x-axis and then going down the y-axis to find the place. The aggregates at that point on the space would be the value queried by the user. In this example, again, those "coordinates" have been "flattened" into a composite key . The key may be regarded as an (X,Y) coordinate of arbitrary dimension. That is, it could be a 4-dimensional coordinate in (W, X, Y, Z) space. Rather than maintain a 4-dimensional cube, a composite key and data store are constructed that are the keys to aggregating data for each point in the space.

[0037] A distributed processing environment is used in the OLAP system described herein to perform bulk updates to the OLAP cube data as new data arrive at the system. The distributed

processing environment allows the cube to be updated in parallel by many computer hosts, thereby radically decreasing the time required to integrate new data into the OLAP system. Further, the system may be similarly updated in real-time, as new data become available they are processed into their two-dimensional projections and inserted in the distributed storage mechanism.

[0038] The system is queried using either a specialized query language that is provided with the system or may be provided through adapters to industry standard query languages for OLAP, including query languages such as XMLA. The queries are processed through a scalable tier of Web-based (HTTP) servers running on computer hosts that are separate from the distributed storage system.

[0039] The distributed processing environment and the distributed storage environment are likewise partitioned into their own tiers. Each tier corresponds to both a class of computer hosts that have been tuned to optimally execute their tasks, and to provide security control by separating the source data, which is not accessible from external systems, from the OLAP data that may be accessed by external systems or over the Internet.

[0040] FIG. 2 is a topology overview of the system computers and data stores of the OLAP system. For example, the "Online Data Center 1" of FIG. 2 corresponds to the "Computer 1" of FIG. 1, and the "Online Data Center N" of FIG. 2 corresponds to the Computer N of FIG. 1. The "Offline Data Center" illustrated in FIG. 2 is a datastore that communicates with the online computers and acts as a data depository to contain large databases for use by the online computers. The Offline Data Center integrates new or existing data in a batch processing mode. The batch processing of the Offline Data Center processes large volumes of source data records to aggregate the data and commit the data to the online OLAP database. FIG. 2 also shows that the online data centers operate in a real-time mode in which the online OLAP database is updated as new data records arrive at the system from the Web API, indicated by the arrows from the Web API showing master data updates and user activity to the Hadoop Distributed File System of each respective online data center. For example, as a user clicks on links of a mobile application on a mobile computing platform, the mobile application sends data to the OLAP cube that corresponds to the links. The data being sent represents user identification, user geographic location, user position within the application, and the like. The sent data is aggregated immediately into the OLAP cube. The dotted lines in FIG. 2 between the blocks for Distributed Table Synch in the online data centers and the Offline Data Center indicate data transfers from

the Offline Data Center to the online data centers to ensure synchronization of the respective databases.

[0041] FIG. 3 illustrates the workflow that occurs in the OLAP system, wherein the FIG. 3 configuration provides a system that utilizes offline background data processing. FIG. 3 shows data updates being received in the upper left corner, to be stored into the Offline Distributed File System. The Distributed File System may be offline, as indicated in FIG. 3, if desired, as only the processed data as described in conjunction with FIG. 1 is required to be online and available for the real-time processing of queries. The parsing, partitioning, and permuting operations (such as described above in conjunction with FIG. 1) are performed in the ETL Workflow Configuration block of FIG. 3, which stores the processed and aggregated data into the Offline Distributed Key-Value Store block and the Offline Hadoop Distributed File System block.

[0042] The offline datastores are then subjected to the online data processing, as indicated in the lower right corner of FIG. 3. The online data is then placed in the Online Distributed Key-Value Store, such as indicated in FIG. 1 by the Computer 1 and Computer N blocks. The online datastore is then accessed by a consumer, such as through the read path of FIG. 1.

[0043] Exemplary Hardware System

[0044] The systems and methods described above may be implemented in a number of ways. One such implementation includes computer devices having various electronic components. For example, components of the system in Figure 1 may, individually or collectively, be implemented with devices having one or more Application Specific Integrated Circuits (ASICs) adapted to perform some or all of the applicable functions in hardware. Alternatively, the functions may be performed by one or more other processing units (or cores), on one or more integrated circuits or processors. In other embodiments, other types of integrated circuits may be used (e.g., Structured/Platform ASICs, Field Programmable Gate Arrays (FPGAs), and other Semi-Custom ICs), which may be programmed in any manner known in the art. The functions of each unit may also be implemented, in whole or in part, with instructions embodied in a memory, formatted to be executed by one or more general or application-specific computer processors.

[0045] Figure 4 is a block diagram of a mobile device 500 such as one of the end user devices 108 illustrated in Figure 1. The mobile device 500 includes a network communications interface 502 through which the mobile device communicates with the partner network 110 or directly

with the recommendations system 100 (Figure 1). A processor 504 controls operations of the mobile device. The processor comprises computer processing circuitry and is typically implemented as one or more integrated circuit chips and associated components. The mobile device includes a memory 506, into which user activity, installed applications, and personal data can be stored. A user input component 508 is the mechanism through which a user can provide controls, responses, and data. The user input component can comprise, for example, a keyboard or numeric pad, voice microphone, touch-sensitive display, or other input mechanism for providing user control and data input. A display 510 provides visual (graphic) output display and an audio component 512 provides audible output for the mobile device.

[0046] Figure 5 provides a block diagram of a computer system 600 for implementing certain functions and operations as described herein. In one embodiment, the system 600 may function as the recommendations system 100 shown in Figure 1 or a partner network computer 110, or one of the front-end or back-end computers illustrated in Figure 2 and Figure 3. It should be noted that Figure 5 is meant only to provide a generalized illustration of various components, any or all of which may be utilized as appropriate. Figure 5, therefore, broadly illustrates how individual system elements may be implemented in a relatively separated or relatively more integrated manner.

[0047] The computational system 600 is shown comprising hardware elements that can be electrically coupled via a system bus 626 (or may otherwise be in communication, as appropriate). The hardware elements can include one or more central processor units (CPUs) 602, including without limitation one or more general-purpose processors and/or one or more special-purpose processors (such as communication processing chips, graphics acceleration chips, and/or the like); one or more input devices 604, that can include, without limitation, a mouse, a keyboard, and/or the like; and one or more output devices 606, which can include without limitation a display device, a printer, audio device, and/or the like.

[0048] The computational system 600 may further include (and/or be in communication with) one or more storage devices 608, which can comprise, without limitation, local and/or network accessible storage and/or can include, without limitation, a disk drive, a drive array, an optical storage device, solid-state storage device such as a random access memory ("RAM"), and/or a read-only memory ("ROM"), which can be programmable, flash-updateable, and/or the like. The computational system 600 might also include a communications subsystem 614, which can include without limitation a modem, a network card (wireless or wired), an infra-red

communication device, a wireless communication device and/or chipset (such as a Bluetooth device, an 802.11 device, a WiFi device, a WiMax device, cellular communication facilities, etc.), and/or the like. The communications subsystem 614 may permit data to be exchanged with a network 615, and/or any other devices described herein. The network 615 may comprise a local area network (LAN) or a network such as the Internet, or a combination. In many embodiments, the computational system 600 will further include a working memory 618, which can include a RAM or ROM device, as described above. The system may optionally include processing acceleration 616 to assist with processing, such as arithmetic computations, graphical computations, and the like.

[0049] The computational system 600 also may comprise software elements, shown as being currently located within the working memory 618, including an operating system 624 and/or other code, such as one or more application programs 622, which may comprise computer programs performing tasks and operations described above, and/or may be designed to implement methods in accordance with the invention and/or configure systems in accordance with the invention, as described herein. Merely by way of example, one or more procedures described with respect to the method(s) discussed above might be implemented as code and/or instructions executable by a computer (and/or a processor within a computer). In one embodiment, the data generating and presenting operations are implemented as application programs 622. In the description herein, references to "interface" and "processor" and "application" should be understood as referring to hardware, software, and combinations of the two, either as independent components (hardware, software, and/or both) for each interface, processor, or application, or as integrated components combined with one or more other components.

[0050] A set of these instructions and/or code may be stored on a computer readable storage medium 610b. In some embodiments, the computer readable storage medium 610b may comprise the storage device(s) 608 described above. In other embodiments, the computer readable storage medium 610b might be incorporated within the computer system. In still other embodiments, the computer readable storage medium 610b might be separate from the computer system (i.e., it may be a removable readable medium, such as a compact disc, etc.), and or might be provided in an installation package, such that the storage medium can be used to program a general purpose computer with the instructions/code stored thereon. These instructions might take the form of executable code, which is executable by the computational system 600 and/or might take the form of source and/or installable code, which, upon compilation and/or

installation on the computational system 600 (e.g., using any of a variety of generally available compilers, installation programs, compression/decompression utilities, etc.), then takes the form of executable code. In these embodiments, the computer readable storage medium 610b may be read by a computer readable storage media reader 610a.

[0051] It will be apparent to those skilled in the art that substantial variations may be made in accordance with specific requirements. For example, customized hardware might also be used, and/or particular elements might be implemented in hardware, software (including portable software, such as applets, etc.), or both. Further, connection to other computing devices such as network input/output devices may be employed.

[0052] In one embodiment, the invention employs local and remote computer systems (such as the computational system 600) to perform methods of the invention. According to a set of embodiments, some or all of the procedures of such methods are performed by the computational system 600 in response to the processor 602 executing one or more sequences of one or more instructions (which might be incorporated into the operating system 624 and/or other code, such as an application program 622) contained in the working memory 618. Such instructions may be read into the working memory 618 from another machine-readable medium, such as one or more of the storage device(s) 608 (or 610). Merely by way of example, execution of the sequences of instructions contained in the working memory 618 might cause the processor(s) 602 to perform one or more procedures of the methods described herein.

[0053] The terms "machine readable medium" and "computer readable medium," as used herein, refer to any medium that participates in providing data that causes a machine to operate in a specific fashion. In an embodiment implemented using the computational system 600, various machine-readable media might be involved in providing instructions/code to processor(s) 602 for execution and/or might be used to store and/or carry such instructions/code (e.g., as signals). In many implementations, a computer readable medium is a physical and/or tangible storage medium. Such a medium may take many forms, including but not limited to, volatile and non-volatile media. Non-volatile computer-readable media includes, for example, optical or magnetic disks, such as the storage device(s) (608 or 610). Volatile computer-readable media includes, without limitation, dynamic memory, such as the working memory 618. In some implementation, data may be carried over transmission media. Transmission media includes coaxial cables, copper wire, and fiber optics, including the wires that comprise the bus 626, as well as the various components of the communication subsystem 614 (and/or the media by

which the communications subsystem 614 provides communication with other devices). Hence, transmission media can also take the form of waves (including, without limitation, radio, acoustic, and/or light waves, such as those generated during radio-wave and infra-red data communications).

[0054] Common forms of physical and/or tangible computer readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, an EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read instructions and/or code.

[0055] Various forms of machine-readable media may be involved in carrying one or more sequences of one or more instructions to the processor(s) 602 for execution. Merely by way of example, the instructions may initially be carried on a magnetic disk and/or optical disc of a remote computer. A remote computer might load the instructions into its dynamic memory and send the instructions as signals over a transmission medium to be received and/or executed by the computational system 600. These signals, which might be in the form of electromagnetic signals, acoustic signals, optical signals, and/or the like, are all examples of carrier waves on which instructions can be encoded, in accordance with various embodiments of the invention.

[0056] The communications subsystem 614 (and/or components thereof) generally will receive the signals, and the bus 626 then might carry the signals (and/or the data, instructions, etc. carried by the signals) to the working memory 618, from which the processor(s) 602 retrieves and executes the instructions. The instructions received by the working memory 618 may optionally be stored on a storage device 608 either before or after execution by the processor(s) 602.

[0057] It will be appreciated that many processing capabilities in addition to those described are possible, without departing from the teachings according to the invention. Further, it should be noted that the methods, systems, and devices discussed above are intended merely to be examples. Various embodiments may omit, substitute, or add various procedures or components as appropriate. For example, it should be appreciated that, in alternative embodiments, the methods may be performed in an order different from that described, and that various steps may be added, omitted, or combined. Also, features described with respect to certain embodiments may be combined in various other embodiments. Different aspects and elements of the

embodiments may be combined in a similar manner. Also, it should be emphasized that technology evolves and, thus, many of the elements are examples and should not be interpreted to limit the scope of the invention.

[0058] Specific details are given in the description to provide a thorough understanding of the embodiments. However, it will be understood by one of ordinary skill in the art that the embodiments may be practiced without these specific details. For example, well-known circuits, processes, algorithms, structures, and techniques have been shown without unnecessary detail in order to avoid obscuring the embodiments. Further, the headings provided herein are intended merely to aid in the clarity of the descriptions of various embodiments, and should not be construed as limiting the scope of the invention or the functionality of any part of the invention. For example, certain methods or components may be implemented as part of other methods or components, even though they are described under different headings.

[0059] Also, it is noted that the embodiments may be described as a process which is depicted as a flow diagram or block diagram. Although each may describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be rearranged. A process may have additional steps not included in the figure.

[0060] Those skilled in the art will understand that the embodiments describe above may be used in any combination with other embodiments described herein or in combination with devices outside of this disclosure without departing from the invention. For example, any of the cascade assembly embodiments may be employed with any of the mechanisms to translate the cascade assembly, or with the optional scoop embodiment. Alternatively, any of the cascade assembly embodiments herein may be employed with an alternative translating mechanism outside of this disclosure without departing from the invention.

[0061] Other variations are within the spirit of the present invention. Thus, while the invention is susceptible to various modifications and alternative constructions, certain illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the invention to the specific form or forms disclosed, but on the contrary, the intention is to cover all modifications, alternative constructions, and equivalents falling within the spirit and scope of the invention, as defined in the appended claims.

[0062] The use of the terms "a" and "an" and "the" and similar referents in the context of describing the invention (especially in the context of the following claims) are to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The terms "comprising," "having," "including," and "containing" are to be construed as open-ended terms (i.e., meaning "including, but not limited to,") unless otherwise noted. The term "connected" is to be construed as partly or wholly contained within, attached to, or joined together, even if there is something intervening. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein, and each separate value is incorporated into the specification as if it were individually recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. The use of any and all examples, or exemplary language (e.g., "such as") provided herein, is intended merely to better illuminate embodiments of the invention and does not pose a limitation on the scope of the invention unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the invention.

[0063] Preferred embodiments of this invention are described herein, including the best mode known to the inventors for carrying out the invention. Variations of those preferred embodiments may become apparent to those of ordinary skill in the art upon reading the foregoing description. The inventors expect skilled artisans to employ such variations as appropriate, and the inventors intend for the invention to be practiced otherwise than as specifically described herein. Accordingly, this invention includes all modifications and equivalents of the subject matter recited in the claims appended hereto as permitted by applicable law. Moreover, any combination of the above-described elements in all possible variations thereof is encompassed by the invention unless otherwise indicated herein or otherwise clearly contradicted by context.

[0064] All references, including publications, patent applications, and patents, cited herein are hereby incorporated by reference to the same extent as if each reference were individually and specifically indicated to be incorporated by reference and were set forth in its entirety herein.

WHAT IS CLAIMED IS:

1. A method of online analytical processing (OLAP) of a computer system, the method comprising:

enriching a data record read from a source database of the computer system by a computer processor with one or more enriched data fields comprising data read from a reference database by the computer processor, wherein the data record includes multiple data fields, such that the enriched data record includes one or more of the multiple data fields of the data record and one or more of the enriched data fields;

partitioning the multiple data fields of the enriched data record into a dimension partition and a metric partition;

permuting the dimension partition and the metric partition of the enriched data record into sub-dimension projections that are mapped to a plurality of storage keys, such that each of the storage keys includes one or more placeholder wildcard values and each of the storage keys is stored into a database of the computer system by the computer processor;

aggregating the storage keys of a permuted sub-dimension projections onto a two-dimensional coordinate vector such that, if the computer processor identifies a permuted storage key having metric values for which a corresponding storage key already exists in the database, then the computer processor aggregates the metric values of the identified storage key with the metric values of the corresponding storage key, and if the computer processor does not identify the permuted storage key as having a corresponding storage key that already exists in the database, then the computer processor writes the metric values of the permuted storage key into the database, comprising initial values for the key combination of dimension values.

2. The method as in claim 1, further comprising:

storing a plurality of the aggregated storage keys across two or more host computer processors, such that the stored aggregated storage keys comprise a distributed associative array.

3. The method as in claim 1, further comprising:

storing the aggregated storage keys of the permuted sub-dimension projection on a persistent medium of the computer processor; and

storing the aggregated storage keys of the permuted sub-dimension projection in cache memory of the computer processor.

4. A computer system for online analytical processing (OLAP) of event stream data, the computer system comprising:

a network communications block adapted for communications by the computer system with a computer network;

a computer processor configured to execute program instructions for

enriching a data record read from a source database of the computer system by the computer processor with one or more enriched data fields comprising data read from a reference database by the computer processor, wherein the data record includes multiple data fields, such that the enriched data record includes one or more of the multiple data fields of the data record and one or more of the enriched data fields,

partitioning the multiple data fields of the enriched data record into a dimension partition and a metric partition,

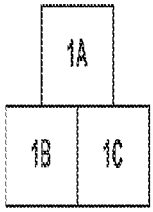
permuting the dimension partition and the metric partition of the enriched data record into sub-dimension projections that are mapped to a plurality of storage keys, such that each of the storage keys includes one or more placeholder wildcard values and each of the storage keys is stored into a database of the computer system by the computer processor, and

aggregating the storage keys of a permuted sub-dimension projections onto a two-dimensional coordinate vector such that, if the computer processor identifies a permuted storage key having metric values for which a corresponding storage key already exists in the database, then the computer processor aggregates the metric values of the identified storage key with the metric values of the corresponding storage key, and if the computer processor does not identify the permuted storage key as having a corresponding storage key that already exists in the database, then the computer processor writes the metric values of the permuted storage key into the database, comprising initial values for the key combination of dimension values.

5. The computer system as in claim 4, wherein the computer processor is further configured to execute program instructions for storing a plurality of the aggregated storage keys across two or more host computer processors, such that the stored aggregated storage keys comprise a distributed associative array.

6. The computer system as in claim 4, wherein the computer processor is further configured to execute program instructions for storing the aggregated storage keys of the

permuted sub-dimension projection on a persistent medium of the computer processor, and for storing the aggregated storage keys of the permuted sub-dimension projection in cache memory of the computer processor.



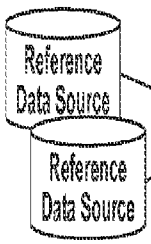




Parsing of Source Data into a Stream of Events

Timestamp	event type	user id	response time
Timestamp	event type	user id	response time
Timestamp	event type	user id	response time
Timestamp	event type	user id	response time

Event Stream



Enrichment of Source Data with Reference Data

Timestamp	event type	user id	user segment	user location	response time	time from last event
Timestamp	event type	user id	user segment	user location	response time	time from last event
Timestamp	event type	user id	user segment	user location	response time	time from last event
Timestamp	event type	user id	user segment	user location	response time	time from last event

Enriched Event Stream

Separation of Events into Dimensions and Metrics

Dimensions					Metrics	
Timestamp	event type	user id	user segment	user location	response time	time from last event
Timestamp	event type	user id	user segment	user location	response time	time from last event
Timestamp	event type	user id	user segment	user location	response time	time from last event
Timestamp	event type	user id	user segment	user location	response time	time from last event

Partitioned Event Stream

Permutation of Sub-Dimension Projections from Partitioned Event Stream

*	*	*	*	*	response time	time from last event
Timestamp	*	*	*	*	response time	time from last event
Timestamp	event type	*	*	*	response time	time from last event
Timestamp	event type	user id	*	*	response time	time from last event
Timestamp	event type	user id	user segment	*	response time	time from last event
Timestamp	event type	user id	user segment	user location	response time	time from last event

Permuted Event Stream

To Fig. 1B To Fig. 1C

FIG. 1A

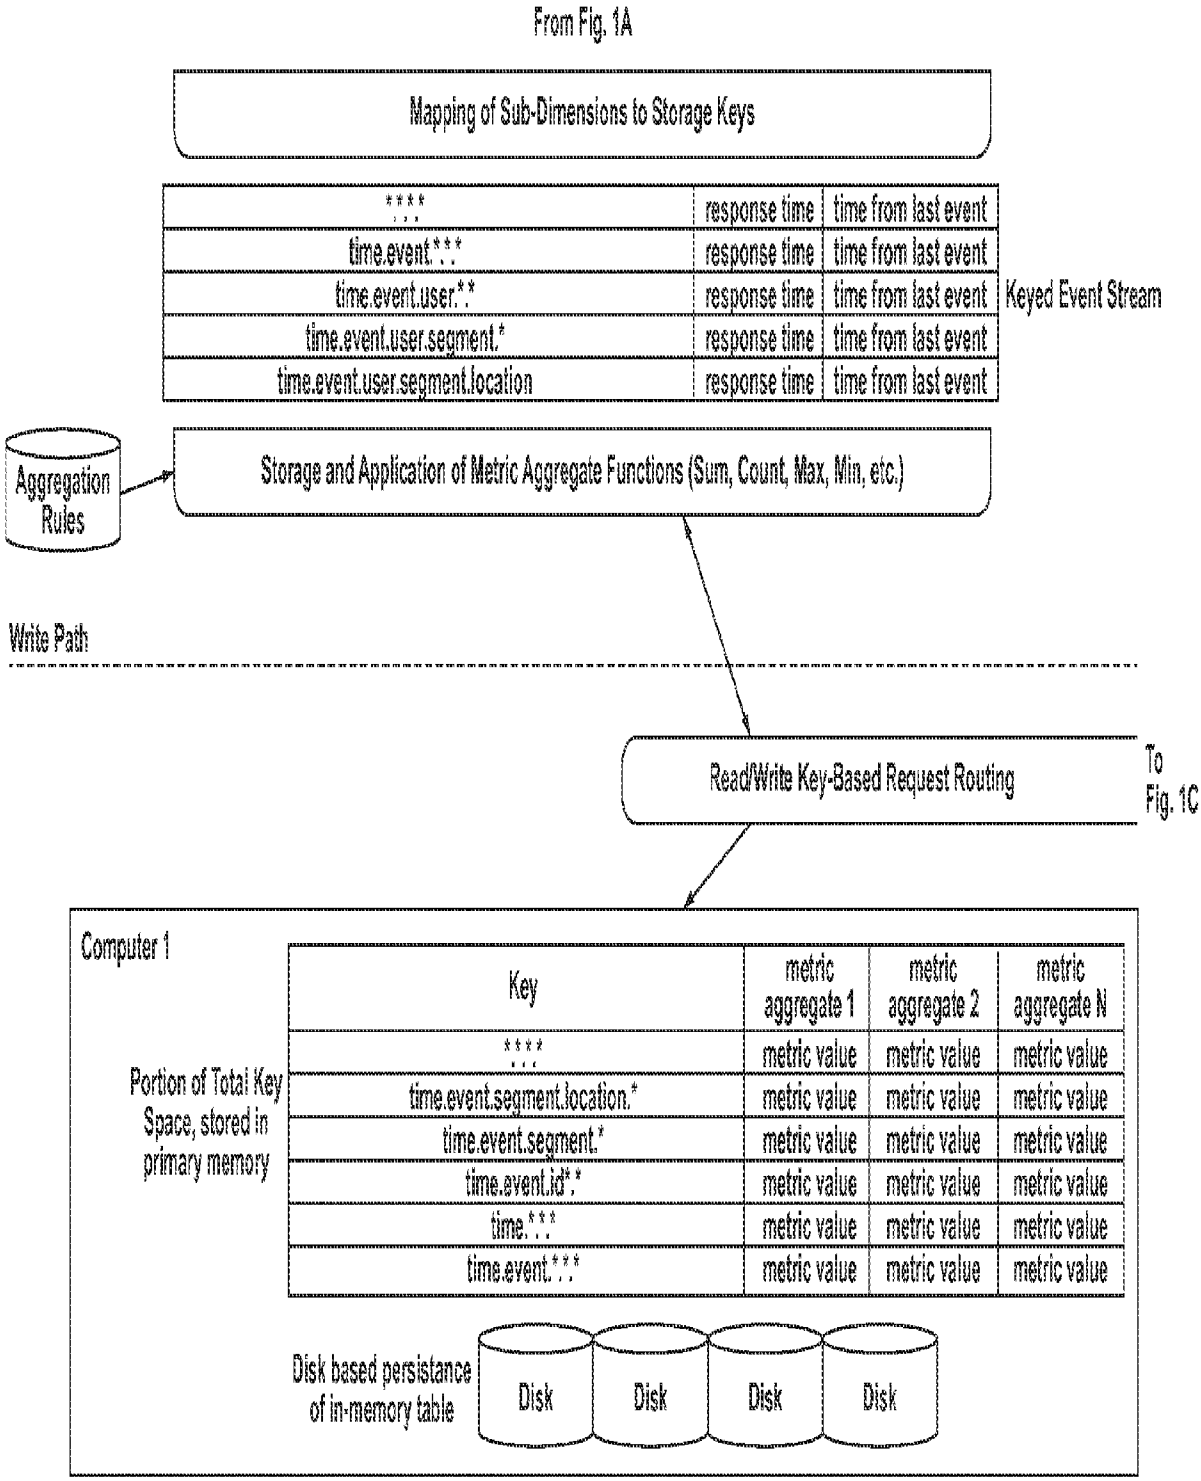

From Fig. 1A
Mapping of Sub-Dimensions to Storage Keys
response time
time from last event
Keyed Event Stream
Aggregation Rules
Storage and Application of Metric Aggregate Functions (Sum, Count, Max, Min, etc.)
Write Path
Read/Write Key-Based Request Routing
To Fig. 1C
Computer 1
Key
metric aggregate 1
metric aggregate 2
metric aggregate N
metric value
Portion of Total Key Space, stored in primary memory
Disk based persistance of in-memory table
Disk
Disk
Disk
Disk

FIG. 1B

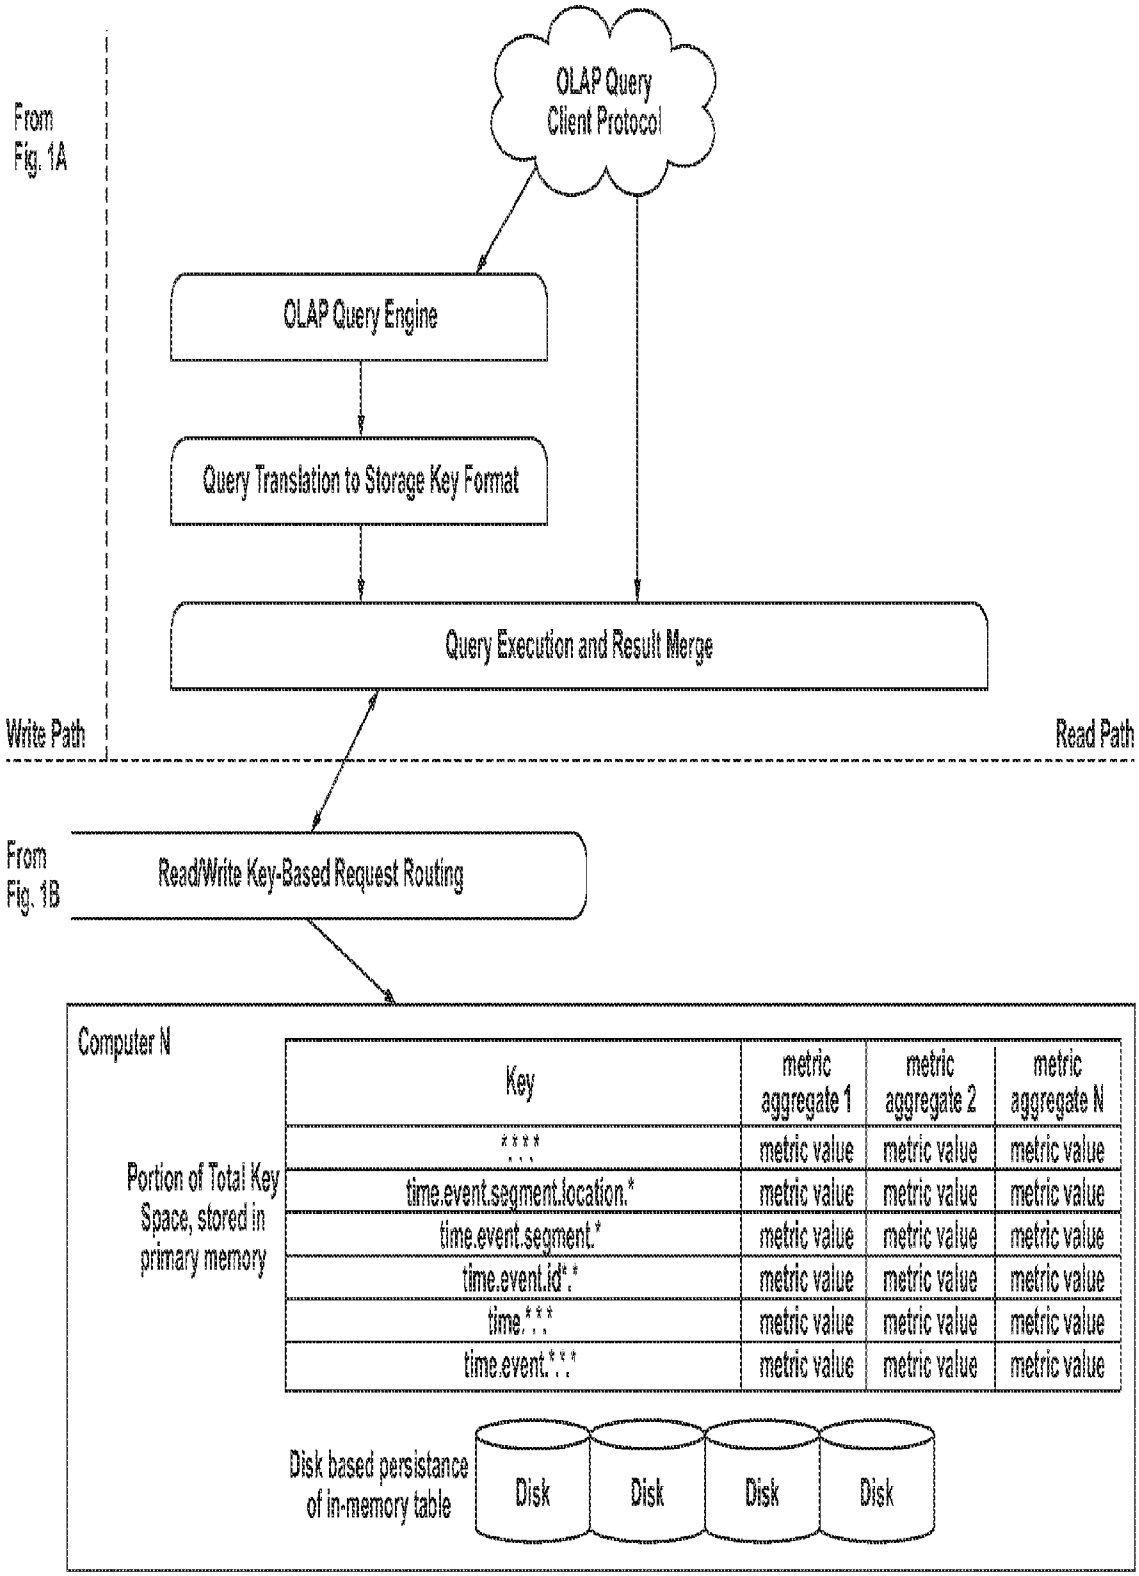

From Fig. 1A
OLAP Query Client Protocol
OLAP Query Engine
Query Translation to Storage Key Format
Query Execution and Result Merge
Write Path
Read Path
From Fig. 1B
Read/Write Key-Based Request Routing
Computer N
Portion of Total Key Space, stored in primary memory
Key
metric aggregate 1
metric aggregate 2
metric aggregate N
****
time.event.segment.location.*
time.event.segment.*
time.event.id*.*
time.***
time.event.***
metric value
Disk based persistance of in-memory table
Disk
Disk
Disk
Disk

FIG. 1C

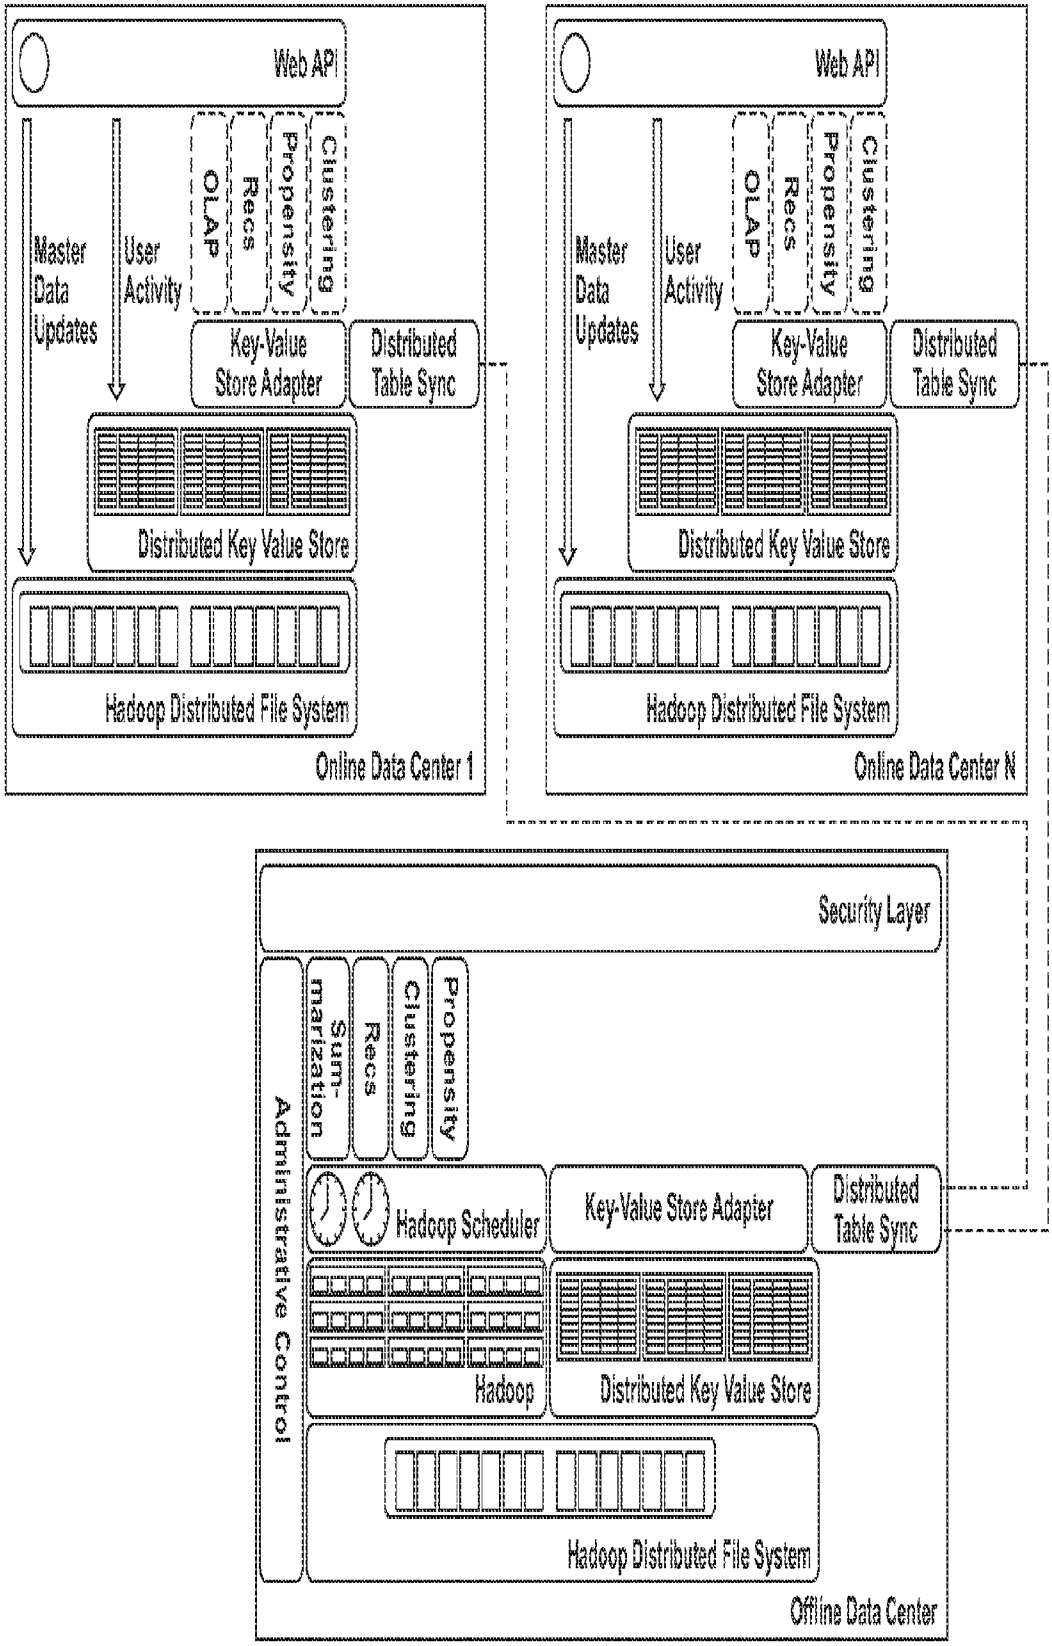

Web API
Master Data Updates
User Activity
OLAP
Recs
Propensity
Clustering
Key-Value Store Adapter
Distributed Table Sync
Distributed Key Value Store
Hadoop Distributed File System
Online Data Center 1
Web API
Master Data Updates
User Activity
OLAP
Recs
Propensity
Clustering
Key-Value Store Adapter
Distributed Table Sync
Distributed Key Value Store
Hadoop Distributed File System
Online Data Center N
Security Layer
Administrative Control
Sum-marization
Recs
Clustering
Propensity
Hadoop Scheduler
Key-Value Store Adapter
Distributed Table Sync
Hadoop
Distributed Key Value Store
Hadoop Distributed File System
Offline Data Center

FIG. 2

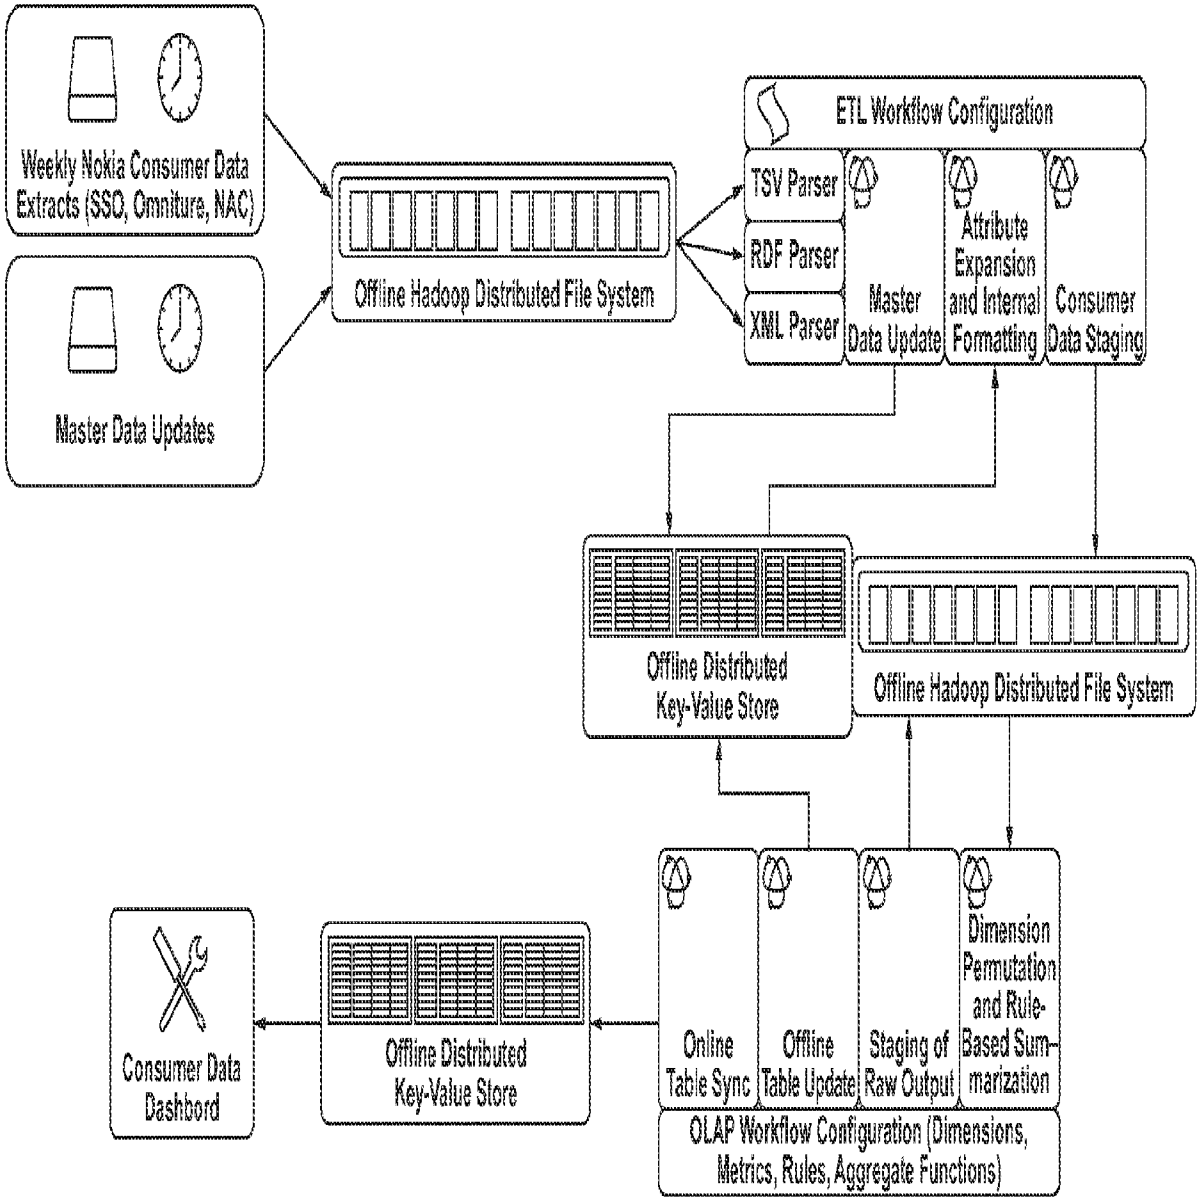

Weekly Nokia Consumer Data Extracts (SSO, Omniture, NAC)
Master Data Updates
Offline Hadoop Distributed File System
ETL Workflow Configuration
TSV Parser
RDF Parser
XML Parser
Master Data Update
Attribute Expansion and Internal Formatting
Consumer Data Staging
Offline Distributed Key-Value Store
Offline Hadoop Distributed File System
Consumer Data Dashbord
Offline Distributed Key-Value Store
Online Table Sync
Offline Table Update
Staging of Raw Output
Dimension Permutation and Rule-Based Sum-marization
OLAP Workflow Configuration (Dimensions, Metrics, Rules, Aggregate Functions)

FIG. 3

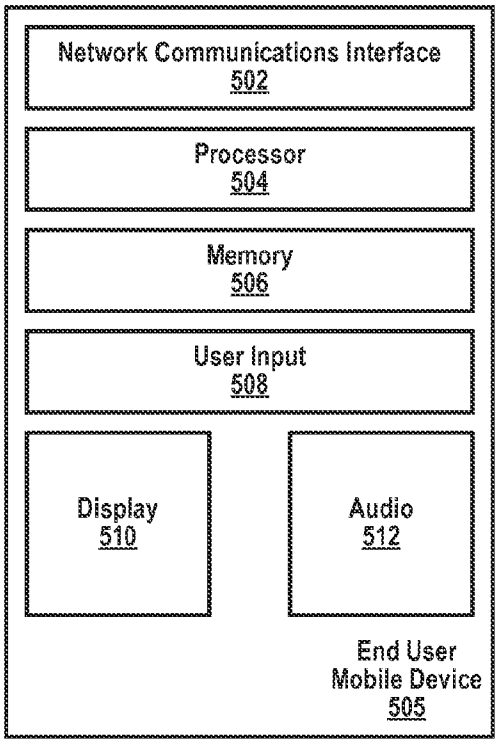

Network Communications Interface
502
Processor
504
Memory
506
User Input
508
Display
510
Audio
512
End User
Mobile Device
505

FIG. 4

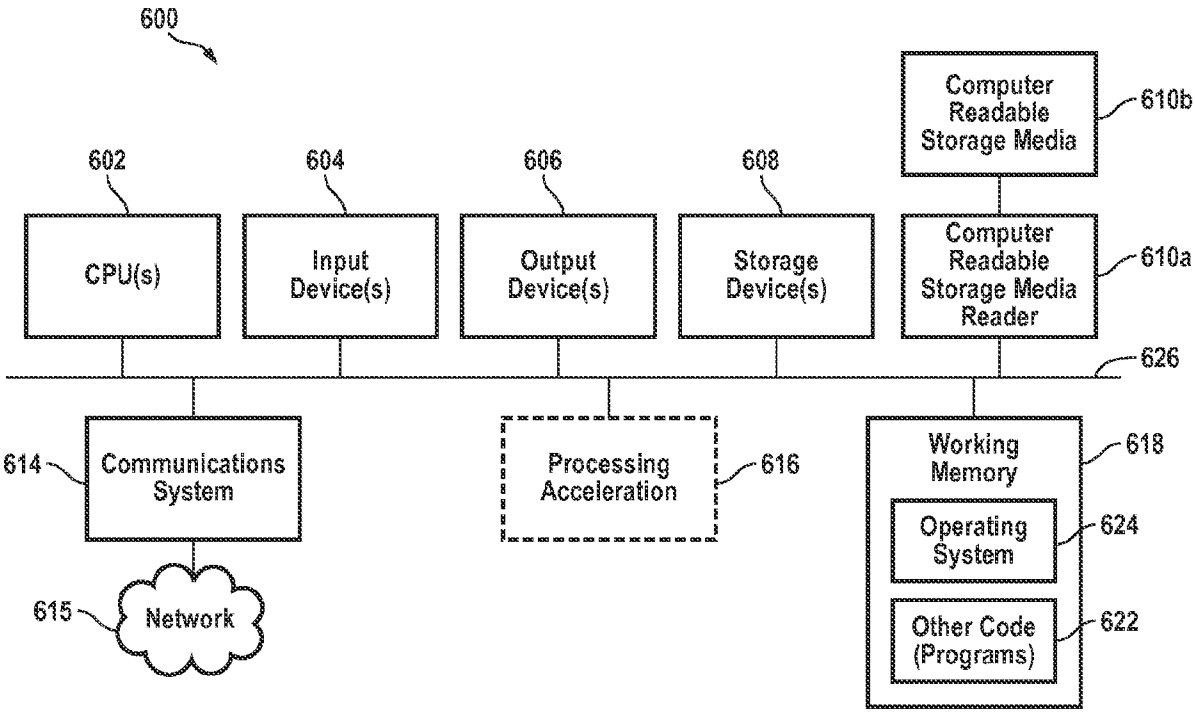

600
Computer Readable Storage Media
610b
602
604
606
608
CPU(s)
Input Device(s)
Output Device(s)
Storage Device(s)
Computer Readable Storage Media Reader
610a
626
614
Communications System
Processing Acceleration
616
Working Memory
618
Operating System
624
615
Network
Other Code (Programs)
622

FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/048006

A. CLASSIFICATION OF SUBJECT MATTER

G06Q 30/02(2012.01)i

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06Q 30/02; G06F 19/00; G06F 17/30; G06Q 10/04; G06Q 10/00

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: online, real time, database, computer, processor, data, record, dimension, metric, wildcard, placeholder, vector

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2010-0250412 A1 (WAGNER) 30 September 2010 See paragraphs [0043]-[0048],[0063]-[0076], and figures 1,9-13.	1-6
A	US 2011-0125668 A1 (COLLIAT et al.) 26 May 2011 See paragraphs [0045]-[0065], and figures 1-6.	1-6
A	KR 10-2009-0059656 A (CHOI SUNG WOOK) 11 June 2009 See abstract, pages 3-5, and figures 1-4.	1-6
A	JP 2007-502466 A (RETEK INC. et al.) 08 February 2007 See abstract, paragraphs [0005]-[0025], and figures 1-6.	1-6
A	US 2011-0161282 A1 (WILLIAMSON) 30 June 2011 See abstract, paragraphs [0015]-[0025], and figures 1-3.	1-6

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents,such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
08 October 2013 (08.10.2013)	**10 October 2013 (10.10.2013)**
Name and mailing address of the ISA/KR Korean Intellectual Property Office 189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701, Republic of Korea Facsimile No. +82-42-472-7140	Authorized officer KIM Sung Gon Telephone No. +82-42-481-8746

Form PCT/ISA/210 (second sheet) (July 2009)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/US2013/048006

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010-0250412 A1	30/09/2010	None	
US 2011-0125668 A1	26/05/2011	US 2011-0040697 A1	17/02/2011
		US 2011-0040698 A1	17/02/2011
		US 2011-0041089 A1	17/02/2011
		US 8510659 B2	13/08/2013
KR 10-2009-0059656 A	11/06/2009	None	
JP 2007-502466 A	08/02/2007	AU 2004-267771 A1	03/03/2005
		AU 2004-267771 B2	16/09/2010
		CA 2534804 A1	03/03/2005
		CN 1867909 A	22/11/2006
		EP 1654673 A2	10/05/2006
		EP 1654673 A4	02/06/2010
		JP 05203607 B2	05/06/2013
		JP 2013-061996 A	04/04/2013
		US 2005-0038768 A1	17/02/2005
		US 7305408 B2	04/12/2007
		WO 2005-019997 A2	03/03/2005
		WO 2005-019997 A3	24/11/2005
US 2011-0161282 A1	30/06/2011	None	