

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
29 November 2007 (29.11.2007)

PCT

(10) International Publication Number
WO 2007/136507 A1

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

(74) Agent: **BINGHAM, Marcel**; Hickman Palermo Truong & Becker LLP, 2055 Gateway Place, Suite 550, San Jose, CA 95110 (US).

(21) International Application Number:
PCT/US2007/010163

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(22) International Filing Date: 24 April 2007 (24.04.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/437,512 18 May 2006 (18.05.2006) US

(71) Applicant (for all designated States except US): **ORACLE INTERNATIONAL CORPORATION** [US/US]; 500 Oracle Parkway, Redwood City, CA 94065 (US).

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(72) Inventors; and

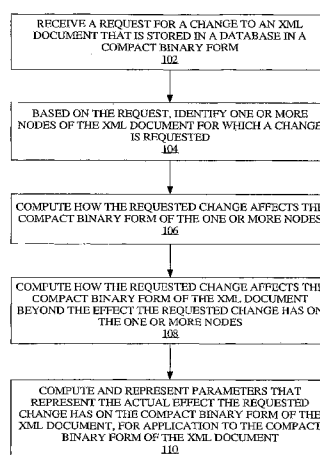
(75) Inventors/Applicants (for US only): **CHANDRASEKAR, Sivansankaran** [IN/US]; 540 Everett Avenue, Palo Alto, California 94301 (US). **GUPTA, Nitin** [IN/US]; 3833 Park Blvd, Apt. 6, Palo Alto, California 94306 (US). **MURTHY, Ravi** [IN/US]; 33227 Jamie Circle, Fremont, California 94555 (US). **AGARWAL, Nipun** [US/US]; 4768 Cheeney Street, Santa Clara, California 95054 (US). **SEDLAR, Eric** [US/US]; 956 Woodside Rd. #3, Redwood City, California 94061 (US).

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

[Continued on next page]

(54) Title: EFFICIENT PIECE-WISE UPDATES OF BINARY ENCODED XML DATA



(57) Abstract: An XML document can be represented in a compact binary form that maintains all of the features of XML data in a useable form. In response to a request for a modification (e.g., insert, delete or update a node) to an XML document that is stored in the compact binary form, a certain representation of the requested modification is computed for application directly to the binary form of the document. Thus, the requested modification is applied directly to the persistently stored binary form without constructing an object tree or materializing the XML document into a corresponding textual form. Taking into account the nature of the binary form in which the document is encoded, the bytes that actually require change are identified, including identifying where in the binary representation the corresponding actual changes need to be made.

WO 2007/136507 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

EFFICIENT PIECE-WISE UPDATES OF BINARY ENCODED XML DATA FIELD OF THE INVENTION

[0001] The present invention relates generally to data management and, more specifically, to techniques for efficiently updating compact binary encoded XML data in a piece-wise manner in a relational database management system (RDBMS).

BACKGROUND

[0002] Use of the Extensible Markup Language (XML) has become a popular and useful technique for representing and exchanging information of any kind, such as exchanging information among computer program applications and services. Consequently, effective and efficient storage and manipulation of XML data has likewise become useful and necessary. Thus, some databases have been augmented to support the storage and manipulation of and access to XML data. One of the primary requirements of applications using XML as their data model is schema flexibility. However, databases for storing XML data are traditionally not optimized for schema flexibility. Although such databases may operate efficiently in a scenario in which a schema is provided and the schema is not prone to change, these databases lack adequate support for XML when the schema is prone to changes or when the schema is loosely structured.

[0003] XML data is self-descriptive (i.e., it contains tags along with data), but the standard XML serialization format is text-based, including the numbers and dates. This results in a significant increase in the size of XML documents compared to other proprietary formats for capturing the same data. The increased size of XML documents causes overhead costs during transmission, due to limited network bandwidths, as well as slower performance of storage and retrieval operations, due to limited disk I/O bandwidth. Hence, a binary encoding form for XML data was introduced which attempts to maximize schema flexibility while still providing storage and querying benefits. This binary encoding form is described in U.S. Patent Appl. No. 11/182,997 filed by Ravi Murthy et al., entitled "Encoding of Hierarchically Organized Data for Efficient Storage and Processing" ("the Murthy application"), the entire content of which is incorporated by reference in its entirety for all purposes as if fully disclosed herein.

[0004] With the encoding format described in the Murthy application, XML data is stored in a compact binary form that maintains all of the features of XML data in a useable form, such as the hierarchical structure underlying the data (e.g., the data model or info set), the notion of elements and attributes, etc. This compact binary format significantly minimizes the overhead due to XML tags. Hence, the encoded XML is more compact than a binary representation of the corresponding textual character

representation. This binary format can be processed more efficiently than parsing because the data is effectively pre-parsed.

[0005] XML documents in a database can be modified by inserting new nodes and by changing or deleting existing nodes, all of which are referred to herein collectively as an update operation. Depending on the nature of the application manipulating XML data, these types of operations can be relatively common. However, existing XML data storage systems do not provide efficient means of updating binary encoded XML documents. Such systems typically load an entire XML document into local memory (e.g., RAM) in the form of an object tree (e.g., a DOM), change the data in memory, and convert the updated DOM tree back into the binary form for storage. This process is generally inefficient and leads to scalability and performance problems because, for one reason, the entire XML document needs to be materialized in local memory.

Additionally, some existing XML data storage systems might provide optimized techniques for updating XML data when there is a very specific XML schema available. However, these systems do not adequately address scenarios in which the XML schema is very unconstrained or scenarios in which there is no XML schema available.

[0006] Hence, based on the foregoing, there is a need for techniques for efficiently updating XML data stored persistently in a database.

[0007] The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0009] FIG. 1 is a flow diagram that illustrates a process for generating a change tuple corresponding to a requested modification to an XML document that is stored persistently in a compact binary form in a database, according to an embodiment of the invention;

[0010] FIG. 2 is a diagram illustrating an example of how an update to compact binary encoded XML data is expressed in a corresponding change tuple, according to an embodiment of the invention; and

[0011] FIG. 3 is a block diagram that illustrates a computer system upon which an embodiment of the invention may be implemented.

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

[0012] In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

FUNCTIONAL OVERVIEW OF EMBODIMENTS

[0013] Techniques for updating XML data stored in a database are described herein. Based on existing technology, an XML document can be represented and interpreted in a compact binary form that maintains all of the features of XML data in a useable form. This compact binary representation of XML may in turn be persistently stored (e.g., on disk) as a large object (e.g., binary large object, or BLOB). In response to a request for a modification (e.g., insert, delete or update a node) to an XML document that is stored in the compact binary form, a certain representation of the requested modification is computed for application directly to the binary form of the XML document. Thus, the requested modification can be applied directly to the persistently stored binary form without constructing an object tree (e.g., a DOM) of the XML document and/or constructing or materializing the XML document into a corresponding textual form.

[0014] The node(s) of the XML document that actually need to be changed are identified. Identifying the affected node(s) includes identifying where in the binary representation the corresponding actual changes need to be made, such as at what byte do the actual changes begin. Thus, at what byte the actual changes begin is included in the representation of the requested modification as a 'destination offset' parameter. However, if using a preferred compact binary form for XML data, identifying where in the binary representation the actual changes begin is not as straightforward as simply identifying the range or ranges of bytes in which the affected nodes or nodes are represented. This is because the preferred binary form takes advantage of known structural properties and constraints in respective XML documents, such as those that are imparted to the documents by way of corresponding XML schemas, for purposes of compacting the XML document generally via tokenization of tags and use of opcodes. For example, if the same tag is repeated multiple times, or if an XML schema constrains the order of elements, then corresponding token IDs can be avoided using the preferred binary form. Thus, computing where in the binary representation the corresponding actual changes need to be made may take into account the nature of the compact binary

form in which the XML document is encoded. Refer to the Murthy application for greater detail regarding the preferred, non-limiting example of a compact binary form of encoding XML data.

[0015] In addition to the destination offset parameter, enough information is included in the representation of the requested modification to be able to interleave the actual changes directly into the binary form of the XML document. For example, certain sets of bytes corresponding to respective updated nodes can be changed, added, or deleted without overwriting the entire XML document.

OPERATING ENVIRONMENT-STORAGE SYSTEM

[0016] A database system typically comprises one or more clients that are communicatively coupled to a database server that is connected to a shared database. "Database server" may refer collectively to a cluster of server instances and machines on which the instances execute. Generally, a server is combination of integrated software components and an allocation of computational resources, such as memory, a node, and processes on the node for executing the integrated software components on a processor, where the combination of the software and computational resources are dedicated to providing a particular type of function on behalf of clients of the server.

[0017] Among other functions of database management, a database server governs and facilitates access to a particular database, processing requests by clients to access the database. In order for a client to interact with a server, a session is established for the client. A session, such as a database session, is a particular connection established from a client to a server, such as a database server. Through a session, the client can issue a series of requests (e.g., requests for data and/or metadata) to the database server.

[0018] A database comprises data and metadata that is stored on a persistent memory mechanism, such as a set of hard disks. Such data and metadata may be stored in a database logically, for example, according to relational and/or object-relational database constructs. Database applications interact with a database server by submitting to the database server commands that cause the database server to perform operations on data stored in a database. A database command that is sent from a database application to a database server contains an original statement of the database command. For the database server to process the commands, the commands must conform to a database language supported by the database server. One non-limiting database language supported by many database servers is known as the Structured Query Language (SQL).

[0019] Generally, data is stored in a database in one or more data containers, each container contains records, and the data within each record is organized into one or more

fields. In relational database systems, the data containers are typically referred to as tables, the records are referred to as rows, and the fields are referred to as columns. In object oriented databases, the data containers are typically referred to as object classes, the records are referred to as objects, and the fields are referred to as attributes. Other database architectures may use other terminology. Systems that implement the present invention are not limited to any particular type of data container or database architecture. However, for the purpose of explanation, the examples and the terminology used herein shall be that typically associated with relational databases. Thus, the terms “table”, “row” and “column” shall be used herein to refer respectively to the data container, record, and field.

[0020] A relational database system may be extended to provide native support for storage, management, and query of (and thus function as a repository for) particular types of data. For example, a traditional relational database system may be augmented with features and technologies to enable storage of XML documents directly in the database, and access to such XML data in either an XML-centric (e.g., using XPath and the XQuery query language) or a relational-centric (e.g., using the SQL query language) manner. Such a repository is at times referred to as an XML repository or an XML database (i.e., “XDB”).

UPDATING XML DOCUMENTS

[0021] XML documents can be modified by inserting, updating, and/ or deleting XML nodes in the documents. A request for a modification to an XML document can be expressed by identifying the XML document, identifying the node(s) to be changed with an XPath or other pathname expression, and providing a new value, if applicable. The set of operations that may be requested are as follows.

(1) Adding new nodes: a user can identify a target node or the parent node by means of an XPath or other pathname expression, and specify the value for the child node to be inserted. The location within an XML document at which a new node is inserted can be identified, for example, by identifying a target node before which the new node is inserted (e.g., insertChild() command) or by identifying a parent node and providing the child node to be inserted as the last child of the identified parent node (e.g., appendChild() command).

(2) Updating existing nodes: a user can identify the target node by means of an XPath or other pathname expression, and specify the new value for the target node to be updated.

(3) Deleting nodes: a user can identify the target node for deletion by means of an XPath or other pathname expression.

[0022] As discussed, traditional XML storage systems typically load the whole XML document into local memory to perform the requested operation. However, this is an inefficient process that results in slow performance. Therefore, according to one embodiment, a stored compact binary representation of the XML document is updated directly based on an XML data modification request. In other words, the XML document may be updated directly on persistent storage (i.e., non-volatile memory, such as disk) without loading the document into local volatile memory. Alternatively, the XML document may be loaded into local volatile memory for updating, but in the compact binary representation rather than as an object tree, where the compact binary representation is more compact than simply representing the textual XML in binary format (see, e.g., the Murthy application).

AN EFFICIENT PROCESS FOR UPDATING XML DOCUMENTS STORED IN A DATABASE

[0023] As mentioned before, a certain representation of a change to an XML document is computed for application directly to a compact binary form of the XML document. According to an embodiment, such a change is represented by a set of parameter values, referred to herein as a "change tuple." Generally, FIG. 1 is a flow diagram depicting a process for generating a change tuple.

[0024] FIG. 1 is a flow diagram that illustrates a process for generating a change tuple corresponding to a requested modification to an XML document that is stored persistently in a compact binary form in a database, according to an embodiment of the invention. The method illustrated in FIG. 1 may be performed by executing one or more sequences of instructions by one or more processors, such as within a computer system (e.g., computer system 300 of FIG. 3).

[0025] At block 102, a request for a modification to an XML document is received, in the context of an XML document that is stored in a database in a compact binary form. For example, a SQL or SQL/XML statement is received at a database server in a database system that supports storage and management of XML data. Such a statement typically identifies (a) an XML document to be updated; (b) one or more pathnames, within the identified XML document, to nodes being updated (e.g., XPaths); and (c) new values for the corresponding nodes being updated.

[0026] At block 104, one or more nodes for which a modification is requested in the XML document are identified based on the modification request. Various approaches to

locating XML data stored in a database may be used to identify the affected nodes based on the pathname specified in the modification request. Non-limiting examples of approaches to locating XML data stored in a database include the use of XML indexes, streaming XPath evaluation (e.g., nondeterministic finite state automation, or NFA), and the like. Furthermore, the compact binary form in which the XML document is stored in persistent storage may ultimately be a binary large object (BLOB). In such a scenario, block 104 may include identifying the exact location within the BLOB at which the affected node or nodes are represented in the serialized binary form.

[0027] At block 106, how the requested modification affects the compact binary form of the one or more nodes is computed. For example, if a node value is changed from “1” to “1000”, how to represent “1000” in the compact binary form is computed. For another example, if a node value is changed from an integer type “1” to a float type “1000”, how to represent the float “1000” in the compact binary form is computed because the compact binary form may store values in their native datatype, as described in the Murthy application. According to one embodiment, the “new value” for the one or more nodes, as represented in the compact binary form, is computed and represented as a parameter in a certain representation (e.g., a change tuple) of the requested modification to the XML document, such as at block 110. According to one embodiment, the new value takes into account any effects the requested modification has on the compact binary form beyond the direct effect the requested modification has on the identified one or more nodes, as described in greater detail hereafter. Hence, the new value represents the net set of bytes inserted into the byte string of the compact binary form to replace the old value.

[0028] Furthermore, a representation of “1000” will likely require more bytes than a representation of “1”, thus, a “new value length” may be computed as part of block 106. According to one embodiment, the new value length for the one or more nodes is computed and represented as a parameter in the change tuple corresponding to the requested modification to the XML document, such as at block 110. According to one embodiment, the new value length takes into account any effects the requested modification has on the compact binary form beyond the direct effect the requested modification has on the identified one or more nodes, as described in greater detail hereafter.

[0029] For an XML document that is stored in a compact binary form (i.e., not a simple binary representation of the character-based XML content), a change to one node may affect the compact binary representation beyond the direct effect on that one node. For example, the Murthy application describes a technique referred to as “array mode

optimization,” whereby if the same tag is repeated multiple times, corresponding token IDs are not repeated. Thus, addition or deletion of a node could affect the tag sequence and, therefore, affect the array mode optimization component of the compact binary encoding scheme. For another example, the Murthy application describes a technique referred to as “schema sequential optimization,” whereby if structure related metadata (e.g., an XML schema) constrains the order of elements to be in a specific order, token IDs can be avoided within the encoding due to exploitation of the XML data model. Thus, a change to an XML schema (e.g., via a DDL operation) could provoke a corresponding change to an XML document in order to remain conforming to the schema. This change to the XML schema and the resultant change to a corresponding XML document could ultimately affect the order of elements within the document and, therefore, the schema sequential optimization component of the compact binary encoding scheme.

[0030] Hence, at block 108, how the requested modification affects the compact binary form of the XML document beyond the effect the requested modification has on the one or more identified nodes is computed. According to one embodiment, this entire net effect on the compact binary form for a particular requested modification is represented as the “new value” for insertion into the compact binary form. Hence, the new value for the entire effect on the compact binary form for a particular requested modification is computed and represented as a parameter in the change tuple corresponding to the requested modification, such as at block 110. Similarly, this entire effect on the compact binary form is represented in the corresponding “new value length” parameter in the change tuple corresponding to the requested modification, such as at block 110. Stated otherwise, the direct effect and the indirect effect, if any, that the requested modification has on the compact binary form are computed and represented together in the corresponding change tuple as corresponding parameters: new value and new value length.

EXPRESSING THE EFFECT OF THE REQUESTED MODIFICATION

[0031] At block 110, parameters that represent the actual effect the requested modification has on the compact binary form of the XML document are computed and represented, such as in a corresponding change tuple. In turn, this change tuple can be applied to the compact binary form of the XML document stored persistently in the database, or locally in volatile memory. According to one embodiment, the effect that the requested modification has on the compact binary form of the XML document is expressed in the form of the following change tuple: <destOffset, changedLen, newValue,

newValueLen>. The names used to refer to the foregoing change tuple parameters may vary from implementation to implementation. Hence, the components of the change tuple are not limited to the foregoing parameter names.

[0032] Note that the nature of a BLOB may vary from implementation to implementation. Thus, some types of BLOBs in which the compact binary representation of the XML document is stored are able to be inserted into and deleted from, whereas other types of BLOBs are not. Stated otherwise, some BLOBs allow for inserting or deleting a group of contiguous bytes anywhere within the BLOB without necessarily having to rewrite the remainder of bytes to the end. By contrast, some BLOBs do not allow inserting or deleting a group of contiguous bytes without overwriting the entire BLOB after the starting insertion/deletion point. Regardless of the nature of the BLOB, the techniques described herein can account for the characteristics of the BLOB when computing change tuple parameters.

[0033] FIG. 2 is a diagram illustrating an example of how an update to compact binary encoded XML data is expressed in a corresponding change tuple, according to an embodiment of the invention.

[0034] The parameter "destOffset" (destination offset) is used to express the point, in the compact binary form, at which the effect of the requested modification begins. Thus, the destination offset identifies, for at least the first atomic node (as appearing in the binary form) being updated via the requested modification, the offset location within the compact binary form at which the requested modification begins. According to one embodiment, the destination offset is represented as a number of bytes from the beginning of the compact binary form. As discussed, the destination offset may not be the actual starting location of an identified node directly affected by the requested modification. Rather, the destination offset identifies the beginning location of the actual change to the compact binary form, with consideration to the nature of the compact binary form. For example, the actual effect of the requested modification on the binary form may begin at a location before the location of the node value in the binary form.

[0035] For the example shown in FIG. 2, assume that only one node is being updated and that the old value of that node is represented in the compact binary form with sixteen bytes, i.e., from bytes 227-243. Assume further that the tag name for the updated node is changed, resulting in a change to the associated array mode optimization component of the compact binary encoding scheme, e.g., a series of five consecutive <lineItem> tags is disrupted by changing the fifth <lineItem> tag to some other tag at the same location in the corresponding XML hierarchy. Therefore, the change to this node has an impact on

the compact binary form beyond just the node being changed, i.e., some bytes before the node need to be changed to reflect the change in array mode optimization. Because of this effect on other bytes, the net effect on the binary form actually begins at byte 222 rather than byte 227. Thus, the destination offset for this update is 222, which indicates that the change to the binary form begins at byte 222 of the byte string.

[0036] The parameter “changedLen” (changed length) is used to express the length of the changed, or old, value of the updated node in its compact binary form. It is noted above that in this example the old value of that node is represented in the compact binary form with sixteen bytes. Thus, the changed length for this update is 16 bytes.

[0037] The parameter “newValue” (new value) is used to express the new value of the entire change to the compact binary form because of the node update, where the new value is expressed in its corresponding compact binary form and includes any effect on the binary form beyond the updated node. The parameter “newValueLen” (new value length) is used to express the length in bytes of the new value of the entire change to the compact binary form because of the node update, where the new value is expressed in its corresponding compact binary form. Thus, the new value length for this update is 21 bytes, i.e., from destination offset byte 222 to the end of the node at byte 243. In this example, a string of bytes representing the old node value, which was expressed with 16 bytes, is replaced with a new string of 21 bytes representing the new node value and this node value’s effect on the rest of the binary form.

[0038] A single user update operation might result in multiple change tuples based on the number of nodes being identified by the XPath in the modification request.

According to one embodiment, these sets of change tuples are represented in an XML change document referred to herein as “atom-diff-XML,” where the “atom” here refers to a primitive change operation. As discussed, the “atom-diff-XML”, which represents the changes to the compact binary form resulting from a modification request, can be applied to the base storage based on the corresponding storage capabilities. For example, if the storage supports piece-wise updates, the primitive change operations represented in the atom-diff-XML are translated directly to the underlying base storage construct. If the storage only supports overwriting capabilities (without the ability to increase/decrease lengths), the overall requested changes represented in the atom-diff-XML are applied to the underlying base storage construct by overwriting everything past the first change (based on the document order).

[0039] In addition to a node update operation, as with the example described in reference to FIG. 2, a node may be inserted or deleted. Inserting a new node into an

XML document adds bytes to the compact binary form. The destination offset identifies where in the byte string the new node is inserted, taking into consideration the new node's effect, if any, on surrounding nodes and the surrounding nodes' corresponding binary representation. Similarly, the new value and new value length corresponding to the node insert operation take into account the new node and the new node's entire effect on the compact binary representation. Even though a new node is being added, there may be a corresponding value for the changed length parameter (the 'old' length) because the new node may affect the compact binary representation of surrounding nodes. Thus, the changed length parameter in such a scenario would represent the old representation of any surrounding nodes that are affected by the insertion of the new node, if any. Likewise, deleting a node from an XML document may replace x number of bytes with y number of bytes, because deleting the node may have an effect on the binary representation of other nodes, whose new binary representation would consist of the y bytes. In summary, whether updating, deleting, or adding a node to an XML document, the values for the change tuple reflect the change to the node as well as the effect the node change has on the compact binary representation of other nodes, if any.

[0040] Hence, techniques are described for applying user update statements without materializing the entire target XML document in memory as a DOM tree, thereby improving the performance and scalability in comparison with other techniques.

HARDWARE OVERVIEW

[0041] FIG. 3 is a block diagram that illustrates a computer system 300 upon which an embodiment of the invention may be implemented. Computer system 300 includes a bus 302 or other communication mechanism for communicating information, and a processor 304 coupled with bus 302 for processing information. Computer system 300 also includes a main memory 306, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 302 for storing information and instructions to be executed by processor 304. Main memory 306 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 304. Computer system 300 further includes a read only memory (ROM) 308 or other static storage device coupled to bus 302 for storing static information and instructions for processor 304. A storage device 310, such as a magnetic disk or optical disk, is provided and coupled to bus 302 for storing information and instructions.

[0042] Computer system 300 may be coupled via bus 302 to a display 312, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 314, including alphanumeric and other keys, is coupled to bus 302 for communicating

information and command selections to processor 304. Another type of user input device is cursor control 316, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 304 and for controlling cursor movement on display 312. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0043] The invention is related to the use of computer system 300 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 300 in response to processor 304 executing one or more sequences of one or more instructions contained in main memory 306. Such instructions may be read into main memory 306 from another machine-readable medium, such as storage device 310. Execution of the sequences of instructions contained in main memory 306 causes processor 304 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0044] The term “machine-readable medium” as used herein refers to any medium that participates in providing data that causes a machine to operation in a specific fashion. In an embodiment implemented using computer system 300, various machine-readable media are involved, for example, in providing instructions to processor 304 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 310. Volatile media includes dynamic memory, such as main memory 306. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 302. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0045] Common forms of machine-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, and 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.

[0046] Various forms of machine-readable media may be involved in carrying one or more sequences of one or more instructions to processor 304 for execution. For example,

the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 300 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 302. Bus 302 carries the data to main memory 306, from which processor 304 retrieves and executes the instructions. The instructions received by main memory 306 may optionally be stored on storage device 310 either before or after execution by processor 304.

[0047] Computer system 300 also includes a communication interface 318 coupled to bus 302. Communication interface 318 provides a two-way data communication coupling to a network link 320 that is connected to a local network 322. For example, communication interface 318 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 318 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 318 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0048] Network link 320 typically provides data communication through one or more networks to other data devices. For example, network link 320 may provide a connection through local network 322 to a host computer 324 or to data equipment operated by an Internet Service Provider (ISP) 326. ISP 326 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 328. Local network 322 and Internet 328 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 320 and through communication interface 318, which carry the digital data to and from computer system 300, are exemplary forms of carrier waves transporting the information.

[0049] Computer system 300 can send messages and receive data, including program code, through the network(s), network link 320 and communication interface 318. In the Internet example, a server 330 might transmit a requested code for an application program through Internet 328, ISP 326, local network 322 and communication interface 318.

[0050] The received code may be executed by processor 304 as it is received, and/or stored in storage device 310, or other non-volatile storage for later execution. In this manner, computer system 300 may obtain application code in the form of a carrier wave.

[0051] In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this application, in the specific form in which such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

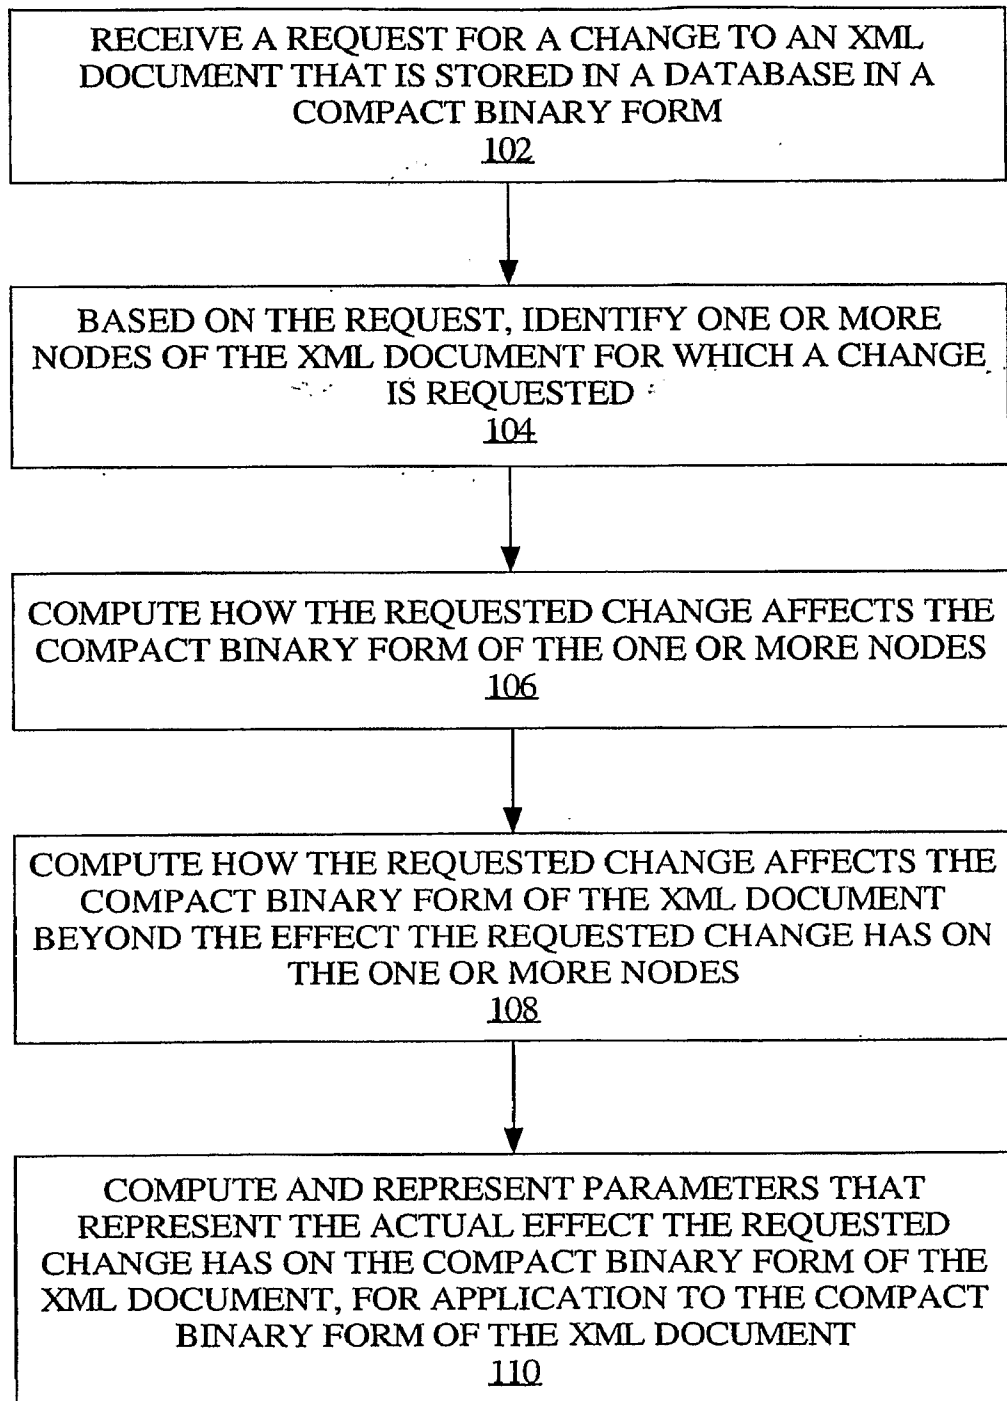
CLAIMS

What is claimed is:

- 1 1. A method comprising performing a machine-executed operation involving
2 instructions for updating XML data stored in a database, wherein the machine-
3 executed operation is at least one of:
4 A) sending the instructions over transmission media;
5 B) receiving the instructions over transmission media;
6 C) storing the instructions onto a machine-readable storage medium; and
7 D) executing the instructions;
8 wherein the instructions are instructions which, when executed by one or more
9 processors, cause the one or more processors to perform the steps of:
10 receiving a request for a modification to an XML document that is stored in a
11 database in a compact binary form;
12 computing parameters that characterize a change to the compact binary form of
13 the XML document, wherein said change reflects said modification; and
14 making the requested modification to the compact binary form of the XML
15 document based on the parameters.
- 1 2. The method of Claim 1, wherein the instructions are instructions which, when
2 executed by one or more processors, cause the one or more processors to perform
3 the steps of:
4 based on the request, identifying one or more nodes of the XML document for
5 which a modification is requested, including identifying a beginning
6 location within the compact binary form at which a first change is to be
7 applied.
- 1 3. The method of Claim 2, wherein computing the parameters comprises computing:
2 a destination offset parameter that represents the beginning location, wherein the
3 beginning location is identified as a number of bytes offset from the
4 beginning of the compact binary form of the XML document.
- 1 4. The method of Claim 2, wherein computing the parameters comprises computing
2 a single XML document that characterizes the modifications to all of the one or
3 more nodes.

- 1 5. The method of Claim 2, wherein the modification to the XML document includes
2 at least one node insert operation or node update operation, and wherein
3 computing the parameters comprises:
4 a destination offset parameter that specifies the beginning location, wherein the
5 beginning location is identified as a number of bytes offset from the
6 beginning of the compact binary form of the XML document;
7 a changed length parameter that specifies a number of bytes of the compact binary
8 form of the XML document used to represent an old binary value
9 associated with a node that is a subject of the node insert operation or node
10 update operation, wherein the old binary value is what is replaced as a
11 result of the node insert operation or node update operation;
12 a new value parameter that represents a new binary value that results from the
13 node insert operation or node update operation, wherein the new binary
14 value is what replaces the old binary value in the compact binary form of
15 the XML document; and
16 a new value length parameter that specifies a number of bytes to represent the new
17 value of the part of the compact binary form of the XML document.
- 1 6. The method of Claim 5, wherein making the change to the compact binary form of
2 the XML document comprises:
3 replacing the old value with the new value beginning at the beginning location.
- 1 7. The method of Claim 1, wherein computing the parameters comprises:
2 computing how the requested modification affects the compact binary form of the
3 XML document beyond an effect the requested modification has on bytes
4 representing one or more nodes from the XML document that are
5 identified for updating based on a pathname expression in the request.
- 1 8. The method of Claim 1, wherein making the requested modification to the
2 compact binary form of the XML document comprises making the requested
3 modification directly to the compact binary form stored on persistent storage.
- 1 9. The method of Claim 8, wherein making the requested modification to the
2 compact binary form of the XML document comprises making the requested
3 modification without constructing an object tree of the XML document.

1/3

**FIG. 1**




FIG. 2




FIG. 3

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2007/010163

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06F17/30

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 practical, search terms used)

EPO-Internal, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	WO 03/107576 A (CERISENT CORP [US]) 24 December 2003 (2003-12-24) pages 3,4, paragraphs 12,13 pages 8,9, paragraph 46-52	1-9
Y	US 2006/021246 A1 (SCHULZE JAMES M [US] ET AL) 2 February 2006 (2006-02-02) cited in the application abstract	1-9
A	WO 2006/026534 A (YAHOO INC [US]; NARSUDE CHETAN [US]) 9 March 2006 (2006-03-09) page 2, paragraph 5 pages 3,4, paragraphs 15,19	1-9
A	US 6 598 055 B1 (KEESEY JAMES LOUIS [US] ET AL) 22 July 2003 (2003-07-22) column 3, line 50 - column 4, line 60	1-9
-/--		

☒ 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 document 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.
- *Z* document member of the same patent family

Date of the actual completion of the international search

5 October 2007

Date of mailing of the international search report

12/10/2007

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

CORREIA MARTINS, F

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2007/010163

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 6 018 747 A (BURNS RANDAL CHILTON [US] ET AL) 25 January 2000 (2000-01-25) abstract -----	1-9

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2007/010163

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 03107576	A	24-12-2003	AU 2003276815 A1	31-12-2003
US 2006021246	A1	02-02-2006	BR 0317044 A	16-11-2005
			CA 2523243 A1	24-06-2004
			EP 1588141 A2	26-10-2005
			KR 20050089812 A	08-09-2005
			WO 2004052074 A2	24-06-2004
			US 2006026915 A1	09-02-2006
			US 2004110460 A1	10-06-2004
WO 2006026534	A	09-03-2006	US 2006059184 A1	16-03-2006
US 6598055	B1	22-07-2003	NONE	
US 6018747	A	25-01-2000	NONE	