

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
24 December 2003 (24.12.2003)


PCT

(10) International Publication Number  
WO 03/107222 A1

- (51) International Patent Classification<sup>7</sup>: **G06F 17/30**
- (21) International Application Number: PCT/US03/18873
- (22) International Filing Date: 13 June 2003 (13.06.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
60/389,066 13 June 2002 (13.06.2002) US
- (71) Applicant: **CERISENT CORPORATION** [US/US];  
2000 Aladeda de las Pulgas, Suite 100, San Mateo, CA  
94403-1269 (US).
- (72) Inventors: **LINDBLAD, Christopher**; 35 Live Oak  
Road, Berkeley, CA 94705 (US). **PEDERSON, Paul**;  
1788 Oak Creek Drive, Apt. 310, Palo Alto, CA 94304  
(US).
- (74) Agents: **YOUNG, Brian, N.** et al.; TOWNSEND AND  
TOWNSEND AND CREW LLP, Two Embarcadero Cen-  
ter, 8th Floor, San Francisco, CA 94111 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT (uti-  
lity model), AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA,  
CH, CN, CO, CR, CU, CZ (utility model), CZ, DE (uti-  
lity model), DE, DK (utility model), DK, DM, DZ, EC, EE  
(utility model), EE, ES, FI (utility model), FI, GB, GD, GE,  
GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ,  
LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN,  
MW, MX, MZ, NI, NO, NZ, OM, PH, PL, PT, RO, RU,  
SC, SD, SE, SG, SK (utility model), SK, SL, TJ, TM, TN,  
TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM,  
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),  
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),  
European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,  
ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO,

[Continued on next page]

(54) Title: PARENT-CHILD QUERY INDEXING FOR XML DATABASES



(57) Abstract: A method for processing queries for a document of elements is provided. The document includes a plurality of subsections where each subsection includes at least a portion of elements in the document. The method comprises: receiving a query for a path of elements in the document of elements; determining a plurality of step queries from the query, each step query including at least a part of the path of elements; for each step query in the plurality of step queries, determining one or more subsections that include elements that correspond to a step query; and determining at least one subsection that includes the path of elements of the query. A result for the query is generated using the at least one subsection.

WO 03/107222 A1



SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

*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.*

**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*

## PARENT-CHILD QUERY INDEXING FOR XML DATABASES

### CROSS-REFERENCES TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 60/389,066,  
5 filed June 13, 2002, entitled "PARENT-CHILD QUERY INDEXING FOR XML  
DATABASES," which disclosure is incorporated herein by reference for all purposes.

The present disclosure is related to the following commonly assigned co pending U.S. Patent  
Applications:

No. \_\_\_\_\_ (Attorney Docket No. 021512 000110US, filed on the same date as  
10 the present application, entitled "A SUBTREE STRUCTURED XML DATABASE"  
(hereinafter "Lindblad I-A");

No. \_\_\_\_\_ (Attorney Docket No. 021512 000310US, filed on the same date as  
the present application, entitled "XML DB TRANSACTIONAL UPDATE SYSTEM"  
(hereinafter "Lindblad III-A"); and

15 No. \_\_\_\_\_ (Attorney Docket No. 021512 000410US, filed on the same date as  
the present application, entitled "XML DATABASE MIXED STRUCTURAL-TEXTUAL  
CLASSIFICATION SYSTEM" (hereinafter "Lindblad IV-A");

The respective disclosures of these applications are incorporated herein by reference for all  
purposes.

### 20 BACKGROUND OF THE INVENTION

[0002] The present invention generally relates to accessing data and more specifically to  
retrieving elements of documents using step queries generated from a query.

[0003] Many languages, such as Extensible Markup Language (XML), define rules that are  
used for structuring data. An XML document is created using the rules to structure data and  
25 includes two parts: the marked up document and the document schema. The marked up part  
of the document encodes a description of the document's storage layout and logical structure.  
The schema part specifies constraints that define XML document structures.

[0004] XML documents are made up of storage units called elements, which may be nested  
to form a hierarchical structure. An element may contain either parsed or unparsed data.  
30 Parsed data is made up of characters, some of which form character data, and some of which  
form the markup; unparsed data is data in its native format. Also, XML elements may have

associated attributes, which may be referred to as name-value pairs. Elements and attributes are described in XML schema where the schema includes, for each element that may occur in the document, a name, the type, the set of attributes, and the set of allowable constituent elements. The relations are represented in a graph with one vertex for each element name, and one edge from an element to each possible constituent.

[0005] In managing XML documents, retrieving elements in the documents for reading or reformatting is often necessary. Accordingly, several query languages have been proposed for searching for and retrieving elements in the XML documents. For example, XQuery, a language derived from an XML query language Quilt and borrowing features from other languages, including XPath, is used for accessing elements in an XML document. XQuery accesses an element using a feature from XPath called an XPath location path expression, which specifies a pattern of elements within the XML document. For example, a query may be of the form A/B/C/D, and is interpreted to specify a pattern of the elements A, B, C, and D within the structure of the XML document. In order to find the desired element D, a system traces the hierarchy of the XML document. The system finds in order, all instances of the element A, all instances of the element B related to element A, all instances of the element C related to the A/B group, and all instances of the element D related to the A/B/C group. Thus, the system processes the XQuery command sequentially, starting from the first element and then to each subsequent element. This method of accessing elements in an XML document becomes time consuming and requires extensive computing power, especially when an element is deeply nested in a hierarchical XML document or a query includes a long path of elements.

#### BRIEF SUMMARY OF THE INVENTION

[0006] In one embodiment of the present invention, a method for processing queries for a document of elements is provided. The document includes a plurality of subsections where each subsection includes at least a portion of elements in the document. The method comprises: receiving a query for a path of elements in the document of elements; determining a plurality of step queries from the query, each step query including at least a part of the path of elements; for each step query in the plurality of step queries, determining one or more subsections that include elements that correspond to a step query; and determining at least one subsection that includes the path of elements of the query. In one embodiment, a result for the query is generated using the at least one subsection.

[0007] A further understanding of the nature and advantages of the invention herein may be realized by reference of the remaining portions in the specifications and the attached drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

- 5 [0008] Fig. 1 is an illustration of a simple XML document including text and markup.
- [0009] Fig. 2 is a schematic representation of the XML document shown in Fig. 1; Fig. 2A illustrates a complete representation the XML document and Fig. 2B illustrates a subtree of the XML document.
- [0010] Fig. 3 is a schematic representation of a more concise XML document.
- 10 [0011] Fig. 4 illustrates a portion of an XML document that includes tags with attributes; Fig. 4A shows the portion in XML format; Fig. 4B is a schematic representation of that portion in graphical form.
- [0012] Fig. 5 shows a more complex example of an XML document, having attributes and varying levels.
- 15 [0013] Fig. 6 is a schematic representation of the XML document shown in Fig. 5, omitting data nodes.
- [0014] Fig. 7 illustrates a possible decomposition of the XML document illustrated in Figs. 5-6.
- [0015] Fig. 8 illustrates the decomposition of Fig. 7 with the addition of link nodes.
- 20 [0016] Fig. 9 illustrates an XQuery server (XQE) according to one embodiment;
- [0017] Fig. 10 illustrates a flow chart for a process for generating database according to one embodiment; and
- [0018] Fig. 11 is a flow chart of a process for generating a result for query according to one embodiment.
- 25 [0019] Fig. 12 depicts a PostingList that may be stored using the structure shown in Fig. 10 according to one embodiment of the present invention.
- [0020] Fig. 13 depicts a PostingList with corresponding scores for each subtree ID according to one embodiment of the present invention.

[0021] Figs. 14A - 14E depict PostingList structures for each subtree according to one embodiment of the present invention.

[0022] Fig. 15A shows a false positive match and Fig. 15B shows a positive match.

#### DETAILED DESCRIPTION OF THE INVENTION

5 [0023] This detailed description illustrates some embodiments of the invention and variations thereof, but should not be taken as a limitation on the scope of the invention. In this description, structured documents are described, along with their processing, storage and use, with XML being the primary example. However, it should be understood that the invention might find applicability in systems other than XML systems, whether they are later-  
10 developed evolutions of XML or entirely different approaches to structuring data.

##### Subtree Storage

[0024] Subtree storage is described in this section, with following sections describing apparatus, methods, structures and the like that might use and store subtrees. Subtree storage is explained with reference to a simple example, but it should be understood that such  
15 techniques are equally applicable to more complex examples.

[0025] Fig. 1 illustrates an XML document 30, including text and markup. Fig. 2A illustrates a schematic representation 32 of XML document 30, wherein schematic representation 12 is shown as a tree (a connected acyclic simple directed graph) with each node of the tree representing an element of the XML document or an element's content,  
20 attribute, the value, etc.

[0026] In a convention used for the figures of the present application, directed edges are oriented from an initial node that is higher on the page than the edge's terminal node, unless otherwise indicated. Nodes are represented by their labels, often with their delimiters. Thus, the root node in Fig. 2A is a "citation" node represented by the label delimited with "<>".  
25 Data nodes are represented by rectangles. In many cases, the data node will be a text string, but other data node types are possible. In many XML files, it is possible to have a tag with no data (e.g., where a sequence such as "<tag></tag>" exists in the XML file). In such cases, the XML file can be represented as shown in Fig. 2A but with some nodes representing tags being leaf nodes in the tree. The present invention is not limited by such variations, so to  
30 focus explanations, the examples here assume that each "tag" node is a parent node to a data node (illustrated by a rectangle) and a tag that does not surround any data is illustrated as a

tag node with an out edge leading to an empty rectangle. Alternatively, the trees could just have leaf nodes that are tag nodes, for tags that do not have any data.

[0027] As used herein, “subtree” refers to a set of nodes with a property that one of the nodes is a root node and all of the other nodes of the set can be reached by following edges in the orientation direction from the root node through zero or more non-root nodes to reach that other node. A subtree might contain one or more overlapping nodes that are also members of other “inner” or “lower” subtrees; nodes beyond a subtree’s overlapping nodes are not generally considered to be part of that subtree. The tree of Fig. 2A could be a subtree, but the subtree of Fig. 2B is more illustrative in that it is a proper subset of the tree illustrated in Fig. 2A.

[0028] To simplify the following description and figures, single letter labels will be used, as in Fig. 3. Note that even with the shorted tags, tree 35 in Fig. 3 represents a document that has essentially the same structure as the document represented by the tree of Fig. 2A.

[0029] Some nodes may contain one or more attributes, which can be expressed as (key, value) pairs associated with nodes. In graph theory terms, the directed edges come in two flavors, one for a parent-child relationship between two tags or between a tag and its data node, and one for linking a tag with an attribute node representing an attribute of that tag. The latter is referred to herein as an “attribute edge”. Thus, adding an attribute (name, value) pair to an XML file would map to adding an attribute edge and an attribute node, followed by an attribute value node to a tree representing that XML file. A tag node can have more than one attribute edge (or zero attribute edges). Attribute nodes have exactly one descendant node, a value node, which is a leaf node and a data node, the value of which is the value from the attribute pair.

[0030] In the tree diagrams used herein, attribute edges sometimes are distinguished from other edges in that the attribute name is indicated with a preceding “@”. Fig. 4A illustrates a portion of XML markup wherein a tag T has an attribute name of “K” and a value of “V”. Fig. 4B illustrates a portion of a tree that is used to represent the XML markup shown in Fig. 4A, including an attribute edge 36, an attribute node 37 and a value node 38. In some instances, tag nodes and attribute nodes are treated the same, such as indexing sequences and the like, but other times are treated differently. To easily distinguish tag nodes and attribute nodes in the illustrated trees, tag nodes are delimited with surrounding angle brackets (“<” and “>”), while attribute nodes are delimited with an initial “@”.

[0031] Fig. 5 et seq. illustrate a more complex example, with multiple levels of tags, some having attributes. Fig. 5 shows a multi-level XML document 40. As is explained later below, Fig. 5 also includes indications 42 of where multi-level XML document 40 might be decomposed into smaller portions. Fig. 6 illustrates a tree 50 that schematically represents multi-level XML document 40 (with a data nodes omitted).

[0032] Fig. 7 shows one decomposition of tree 50 with subtree borders 52 that correspond to indications 42. Each subtree border 52 defines a subtree; each subtree has a subtree root node and zero or more descendant nodes and some of the descendant nodes might in turn be subtree root nodes for lower subtrees. In this example, the decomposition points are entirely determined by tag labels (e.g., each tag with a label “c” becomes a root node for a separate subtree, with the original tree root node being the root node of a subtree extending down to the first instances of tags having tag labels “c”). In other examples, decomposition might be done using a different set of rules. For example, the decomposition rules might be to break at either a “c” tag or an “f” tag, break at a “d” tag when preceded by an “r” tag, etc.

Decomposition rules need not be specific to tag names, but can specify breaks upon occurrence of other conditions, such as reaching a certain size of subtree or subtree content. Some decomposition rules might be parameterized where parameters are supplied by users and/or administrators (e.g., “break whenever a tag is encountered that matches a label the user specifies”, or more generally, when a user-specified regular expression or other condition occurs).

[0033] Note from Fig. 7 that subtrees overlap. In a subtree decomposition process, such as one prior to storing subtrees in a database or processing subtrees, it is often useful to have nonoverlapping subtree borders. Assume that two subtrees overlap as they both include a common node. The subtree that contains the common node and parent(s) of the common node is referred to herein as the upper overlapping subtree, while the subtree that contains the common node and child(ren) of the common node is referred to herein as the lower overlapping subtree.

[0034] Fig. 8 illustrates one approach to having nonoverlapping subtrees, namely by introducing the construct of link nodes 60. For each common node, an upper link node is added to the upper subtree and a lower link node is added to the lower subtree. These link nodes are shown in the figures by squares. The upper link node contains a pointer to the lower link node, which in turn contains a pointer to the root node of the lower overlapping

subtree (which was the common node), while the lower link node contains a pointer to the upper link node, which in turn contains a pointer to the parent node of what was the common node. Each link node might also hold a copy of the other link node's label possibly along with other information. Thus, the upper link node may hold a copy of the lower subtree's root node label and the lower link node may hold a copy of the upper subtree's node label for the parent of what was the common node.

[0035] An XQuery may include an XPath location that indicates a path expression of elements. XPath location path expressions have the form "*name\_a/name\_b/.../name\_z*", and specify a pattern of elements within the XML document hierarchical element structure. The terms *name\_a*, *name\_b*, ... refer to elements or attributes of the XML document. The XPath location path expression is used to specify a desired element that is to be retrieved. For example, if all authors' last names appearing within a "citation" fragment are desired, an XPath location path pattern is *A/C/E* or *citation/author/last*. For discussion purposes, the queries processed will be XQuery and XPath queries for XML documents; however, it will be understood that a person skilled in the art will appreciate other queries that may be processed for other documents. Thus, embodiments of the present invention are not limited to XML documents and XQueries.

[0036] Fig. 9 illustrates an XQuery server (XQE) 200 according to one embodiment. XQE 200 includes a document processor 204 and a query processor 218. Document processor 204 generates step queries and step query results from documents 202 and stores the step queries and step query results in a database 212. In one embodiment, documents 202 are parsed documents. For example, parsed documents are created by an XML parsing process. The parsing process accepts XML textual inputs (serialized XML), analyzes the element structure of these documents, and outputs a data structure that represents the input document as a linked collection of element nodes linked to attribute nodes and child element nodes. The parsed XML document also may contain text nodes, processing instruction nodes, and comment nodes.

### Overview

[0037] Query processor 218 receives a query 219 for elements in documents 202 and generates step queries from a query 219. In one embodiment, query 219 is a parsed query. Parsed queries are created by an XQuery parsing process. The XQuery parsing process accepts XQuery textual inputs, analyzes their grammatical structure, and outputs a data

structure that represents the Xquery query as a linked collection of expression nodes. For example, each query expression of the form 'A *op* B' is represented as an *op-node* with two children nodes representing the subexpressions A and B. The results from the step queries are retrieved from database 212 and a result for query 219 is determined. For example, an  
 5 intersection of the results is taken to generate the result for query 219.

### Document Processing

[0038] XQE 200 receives documents 202, such as the XML document of Fig. 1. Documents 202 are passed to a document processor 204, which includes a step query generator 206, a canonicalizer 208, a hash key generator 210, and a step query result  
 10 generator 214. After receiving documents 202, step query generator 206 generates step queries from documents 202. The step queries are patterns from the hierarchical structure of elements in documents 202. For example, the step queries are relationships between elements that may be part of possible queries for elements in document 202. The number of steps, K, in a query represents a number of levels of relationships between elements. A query  
 15 may be for any number of K steps. A larger K means faster execution but more space is required to store the larger step queries. A smaller K means slower execution but less space is required. For example, a two-step query may be a query for a parent node and its child node and a three-step query may be a query for a parent, its child, and the child's child. In one embodiment, one-step, two-step, three-step, and four-step queries may be generated from  
 20 elements in documents 202. These step queries may take the form of these patterns for:

[0039] one-step queries:

- (a) *elem*,
- (b) *word::wrd*;

two-step queries:

- 25 (c) *elem/word::wrd*,
- (d) *elem/word::[string]*,
- (e) *elem/child*,
- (f) *elem/@attr*;

three-step queries:

- 30 (g) *elem/@attr/word::wrd*,
- (h) *elem/@attr/word::[string]*,
- (i) *elem/child/word::wrd*,

- (j) *elem/child/word::[string]*,
- (k) *parent/elem/child*;

and four-step queries:

- (l) *elem/child/@attr/word::wrđ*,
- 5 (m) *elem/child/@attr/word::[string]*,
- (n) *parent/elem/child/word::wrđ*,
- (o) *parent/elem/child/word::[string]*,
- (p) *grandp/parent/elem/child*.

[0040] It will be understood that step queries are not limited to the above queries and other  
10 step queries may be used and derived from documents 202.

[0041] Once the step queries are generated from documents 202, the step queries are passed  
to canonicalizer 208. Canonicalizer 208 reduces each step query to its canonical form. For  
example, the one-step queries are reduced to the following canonical forms:

- (a) *elem* -> *elem*,
- 15 (b) *word::wrđ* -> *wrd*;

the two-step queries to the following canonical forms:

- (c) *elem/word::wrđ* -> *elem#word("wrđ")*,
- (d) *elem/word::[string]* -> *elem#string*,
- (e) *elem/child* -> *elem/#child*,
- 20 (f) *elem/@attr* -> *elem#/@#attr*;

the three-step queries to the following canonical forms:

- (g) *elem/@attr/word::wrđ* -> *elem#/@#attr#word("wrđ")*,
- (h) *elem/@attr/word::[string]* -> *elem#/@#attr#string*,
- (i) *elem/child/word::wrđ* -> *elem##child#word("wrđ")*,
- 25 (j) *elem/child/word::[string]* -> *elem##child#string*,
- (k) *parent/elem/child* -> *parent##elem##child*;

and the four-step queries to the following canonical forms:

- (l) *elem/child/@attr/word::wrđ* -> *elem##child/@#attr#word("wrđ")*,
- (m) *elem/child/@attr/word::[string]* -> *elem##child/@#attr#string*,
- 30 (n) *parent/elem/child/word::wrđ* -> *parent##elem##child#word("wrđ")*,
- (o) *parent/elem/child/word::[string]* -> *parent##elem##child#string*,
- (p) *grandp/parent/elem/child* -> *grandp##parent##elem##child*;

[0042] Once the step queries are reduced to their canonical form, the step queries are passed to hash key generator 210.

[0043] Hash key generator 210 generates hash keys for each canonical form that may be used for indexing results for each step query. Although hash keys are described, it should be understood that any reference to a storage location may be used. In one embodiment, a 64-bit hash value is computed for each canonical form. Individual names such as parent, element, and child, as well as the literals word (" "), /@, and / generate hash values by direct application of a 64-bit hashing function in hash key generator 210. Also, terms (tokens) separated by the hash mark "#" may be composed by applying either the formula  $A\#B = \text{hash64}(A) * 5 + \text{hash64}(B)$ , or the formula  $A\#B = \text{hash64}(\text{hash64}(A), B)$ , where hash64 represents the hashing function. The latter formula expresses a general compositional mechanism for forming the hash key for two tokens using previously computed hash values for the first token. Hash key generator 210 uses hash value caches and hash composition to compute hash keys for all the indexable step queries.

[0044] For example, the hash key for *elem#@attr* is computed by the hashing function as:

$$\text{hash64}(\text{elem}) * 5 + (\text{hash64}(/@) + \text{hash64}(\text{attr}));$$

and the hash key for *A#B#C* is computed as:

$$\text{hash64}(\text{hash64}(\text{hash64}(A), B), C)$$

and the hash key for *A#B#C#...#Y#Z* is computed as

$$\text{hash64}(\text{hash64}(\dots(\text{hash64}(\text{hash64}(\text{hash64}(A), B), C), \dots, Y), Z).$$

[0045] Additionally, the hash key for a string value, such as: *string = word\_1 word\_2 ... word\_n*, is computed by composition across the word tokens within the string. Thus, the hash key is computed as follows:

$$(\dots(\text{hash64}(\text{word}_1) * 5 + \text{hash64}(\text{word}_2) * 5 + \dots) * 5 + \text{hash64}(\text{word}_n)).$$

[0046] The hash keys generated from hash key generator 210 are stored in an index 213 of database 212. Also, the hash keys for each canonical step query are used to access step query results stored in index 213 of database 212. The step query results are generated by step query result generator 214. Step query result generator 214 receives the step queries generated from step query generator 206 and generates the results for each step query using the hierarchical structure of documents 202 corresponding to the step query. In one embodiment, the step query results may be the element(s) corresponding to the step query. In another embodiment, the step query results for a step query are one or more sub-tree IDs

corresponding to the XML fragment for the step query. The step query and corresponding step query results may be stored as a PostingList, which will be described below.

Additionally, a frequency count of how many times the step query result occurs within the XML fragment is connected with the step query result. Once the results for the step queries are determined, step query result generator 214 stores the results in index 213. In one embodiment, index 213 includes, but is not limited to, the results of all atomic one-step queries of the forms:

(a) find all elements with a given name,

(b) find all elements containing a given word;

10 in addition, it includes the results of all two-step queries of the forms:

(c) find all elements of a given name whose text content contains a given word,

(d) find all elements of a given name whose text content equals a given string,

(e) find all elements of a given name with a child element of a given name,

(f) find all elements of a given name with an attribute of a given name;

15 in addition, it includes the results of all three-step queries of the forms:

(g) find all elements of a given name with an attribute of a given name whose value contains a given word,

(h) find all elements of a given name with an attribute of a given name whose value equals a given string,

20 (i) find all elements of a given name with a child element of a given name whose text content contains a given word,

(j) find all elements of a given name with a child element of a given name whose text content equals a given string,

25 (k) find all elements of a given name with a parent of a given name and a child element of a given name;

and in addition, it includes the results of all three-step queries of the forms:

(l) find all elements of a given name with a child element of a given name with an attribute of a given name whose value contains a given word,

30 (m) find all elements of a given name with a child element of given name with an attribute of a given name whose value equals a given string,

(n) find all elements of a given name with a parent element of a given name with a child element of a given name whose text content contains a given word,

(o) find all elements of a given name with a parent element of a given name with an

element of a given name with a child element of a given name whose text content equals a given string,

- (p) find all elements of a given name with a grandparent element of a given name with a parent element of a given name with an element of a given name and a child element of a given name;

5

**[0047]** It will be understood that the step query results are not limited to the above possibilities and may store atomic query results up to any fixed finite level.

**[0048]** In one embodiment, index 213 is an inverted file index. The inverted file index maps terms to PostingLists. The terms correspond to textual units extracted from a collection of documents 202 or document fragments from documents 202, and PostingLists describe where and how often each term appeared within a given document or document fragment from documents 202. In one embodiment, 'terms' are the atomic text units of document 202. Terms are generated by 'tokenizing' the text content of the document. Text is tokenized through a process of table lookup for each character to determine if that character is a word constituent, white space, or punctuation. Word constituent characters delimited by either spaces or punctuation are accumulated as 'tokens'. Canonicalized step queries are also terms.

10

15

**[0049]** In one embodiment, a hash key is stored in a memory-map list index file in index 213 whose entries contain (key, offset) pairs, where the offset describes the absolute location within a Listdata file where the list of results for the step query may be found. Thus, the Listdata file includes a reference to the step query results. In one embodiment, the step query results are stored as a compressed list of (subtree-id, frequency-count) pairs. A subtree-id uniquely identifies the XML fragment matching the atomic step query, and the frequency-count describes the approximate number of times that the match occurred within document 202 or the document fragment of document 202.

20

**[0050]** In one embodiment, a list of results in the ListData file may be referred to as the PostingList. The PostingList includes the unique subtree-id identifier of the corresponding result of the step query. Additionally, the PostingList includes a score, which is a normalized frequency count. For example, index 213 stores, for each term, at a location determined by the hash key of that term, a PostingList containing references to the subtrees containing the term along with a normalized frequency count (score) that approximates the number of occurrences of the term within the subtree. In one embodiment, the sequence of nodes returned by the function search may be ordered by a 'relevance' score. The relevance of a

30

node to the specified query is a complex function that depends on the frequency the query terms appear in the text of the query nodes, the frequency the query terms appear across the entire database, and the quality score attached to a given node. The quality score is further described in Linblad IV-A. In one embodiment, the PostingLists are stored in a compressed  
5 format. Although the PostingList is described, it will be understood that other lists may be used to store step query results.

[0051] Each hash value provides an index into a memory-mapped ListIndex file of fixed-length records. Each record contains a pair including a hash key and a fixed-width file offset. The file offset describes the location within a secondary ListData heap file where the  
10 PostingLists are stored. Binary search finds the (key, offset) pair within the ListIndex file, then a single random access I/O to the ListData file locates the first block of PostingList data. In most cases one data block contains the entire PostingList. But if not, and the PostingList exceeds the size of one data block, then subsequent sequential I/O's fetch the remainder of the list. The number of I/O is proportional the length of the PostingList divided by the packing  
15 factor -- that is, the number of individual postings per block.

[0052] In one embodiment, the format uses unary-log-log variable length bit encodings for subtree id's and scores. Furthermore, both subtree id's and scores may be kept in a differential form where each Posting stores only the encoded difference from the preceding subtree id and score. Large PostingLists typically have long strings of consecutive subtree  
20 id's with scores that are mostly equal. The PostingList formats encode the consecutive runs using only one or two bits for the delta(id) (the id differential), and delta (score) (the score differential). Large PostingLists are stored with markers containing sufficient information to allow a search process to skip forward across blocks of Postings (a "skip-list" structure). The skip-list block size a configurable parameter.

[0053] For any choice of the skip-list block size parameter, three cases may arise: (1) the PostingList size is less than fifteen, (2) the PostingList size is less than or equal to one block, and (3) the PostingList size exceeds a single block. In the following description, the square brackets [ ] indicate 'unary-log-log variable length bit encoding'. The parentheses indicate bit  
25 fields of a specified size, (e.g., *length*(0:3) means a 4-bit field). The notations {0} and {1} indicate constant bits equal to 0 and 1, respectively.  
30

[0054] In case (1) the length, being less than 15, occupies four leading bits, and the rest of the format is packed with variable-length bit encodings of differential subtree *id*'s and *scores*:

*length*(0:3), [*id*0], [*score*0],  
           [*id*1 - *id*0], [*score*1 - *score*0],  
           [*id*2 - *id*1], [*score*2 - *score*1],  
           [*id*3 - *id*2], [*score*3 - *score*2], ...

5

[0055] In case (2), the four leading bits are all set to 0, and the format is:

{0}(0:3), [*length*],  
           [*id*0], [*score*0],  
           [*id*1 - *id*0], [*score*1 - *score*0],  
           [*id*2 - *id*1], [*score*2 - *score*1],  
           [*id*3 - *id*2], [*score*3 - *score*2], ...

10

15 [0056] And in case (3), the four leading bits are all set to 1 and the format is:

{1}(0:3), [*length*], *Block*0, *Block*1, *Block*2, ...

Each *Block* has the format:

20

*maxSubtreeID*[0:32], *numPostings*[0:15], *numWords*[0:15],  
           [*id*0], [*score*0],  
           [*id*1 - *id*0], [*score*1 - *score*0],  
           [*id*2 - *id*1], [*score*2 - *score*1],  
           [*id*3 - *id*2], [*score*3 - *score*2], ...

25

[0057] *MaxSubtreeID* bounds the ordinal size of any subtree id appearing in the block; *numPostings* bounds the number of Postings in the block and *numWords* is the size of the block in 32-bit words.

[0058] A search for a given subtree *id* proceeds by scanning down the list: if *maxSubtreeID*  
 30 is smaller than the given id, then the process skips forward to the start of the next block by incrementing the list offset by *numWords*.

30

[0059] The granularity of index 213 will now be described. More details relating to index 213 and storage of subtree IDs are disclosed in Linblad I-A. Index 213 stores the SubTree ids. The result of searching database 212 with step queries is a list of SubTrees satisfying the  
 35 step queries. The system synthesizes per-element search query results by loading whole SubTrees into memory of XQE 200 and then seeking within the SubTree for specific elements, attributes, text content, or any of the combinations of elements, attributes and content described above. The SubTree represents a unit of locality. The indexes are designed to speed up queries that can be resolved by locating a contiguous fragment of the  
 40 original XML document and then navigating within that fragment.

40

[0060] In one example, referring to Figs. 1 and 2, document processor 204 may receive the document fragment of Fig. 1 and generate step queries of the form:

one step queries:

A, B, C, D, E, F; and

5  $term\_1, term\_2, term\_3, \dots$  for each term appearing in A, B, C, ...;

two-step queries:

A/B, A/C, A/D, C/E, and C/F, and

A/word( $term\_1$ ), A/word( $term\_2$ ), A/word( $term\_3$ ),...

three-step queries:

10 A/C/E, and A/C/F.

[0061] Each of the above step queries may be assigned a hash value and stored in index 213. The results of the step queries are then computed and stored in step query database 216. The hash value may then be used to look up the step query results, which point to one or more sub-tree IDs for the XML fragment corresponding to the step query. For example, the step query A/B includes the sub-tree ID for the citation/title fragment.

### Query Processing

[0062] One embodiment of query processor 218, which includes an optimizer 220, a step query generator 222, a composer 224, and an intersector 226, will now be described. Query processor 218 receives query 219, generates step queries from query 219, uses the generated step queries to retrieve the pre-computed step query results in database 212, and uses the step queries to output a query result. Thus, when a query is received for documents 202, the results for step queries generated from the query are already known.

[0063] After receiving query 219, query processor 218 sends query 219 to optimizer 220, which may optimize the query if necessary. The optimization process will be described in more detail below. The optimized query is then sent to a step query generator 222.

[0064] Step query generator 222 generates step queries from query 219. As described above, with reference to step query generator 206, step query generator 222 breaks query 219 into step queries, such as one-step queries, two-step queries, three-step queries, and four-step queries. For example, step query generator 222 reduces or decomposes query 219 of a form:

30  $aa\_1/aa\_2/aa\_3/aa\_4/.../aa\_(n-1)/aa\_n$  (where the ellipses indicate that any finite number of additional steps may appear in the query)

to a sequence of two-step queries as follows:

*aa\_1/aa\_2, aa\_2/aa\_3, aa\_3/aa\_4, ..., aa\_(n-2)/aa\_(n-1), aa\_(n-1)/aa\_n.*

5

[0065] Queries containing trailing attribute specifications, as in

*aa\_1/aa\_2/aa\_3/aa\_4/.../@aa\_n*, are reduced to a sequence of two-step queries as follows:

*aa\_1/aa\_2, aa\_2/aa\_3, aa\_3/aa\_4, ..., aa\_(n-2)/aa\_(n-1), aa\_(n-1)/@aa\_n.*

10

[0066] Queries containing trailing word specifications, as in

*aa\_1/aa\_2/aa\_3/.../aa\_n/word:: wrd*, are reduced to a sequence of two-step queries as follows:

*aa\_1/aa\_2, aa\_2/aa\_3, ..., aa\_(n-1)/aa\_n, aa\_n/word:: wrd.*

15

[0067] Queries containing trailing attribute word specifications, as in

*aa\_1/aa\_2/aa\_3/.../aa\_(n-1)/@aa\_n/word:: wrd*, are reduced to a sequence of two-step and three-step queries as follows:

20 *aa\_1/aa\_2, aa\_2/aa\_3, ..., aa\_(n-2)/aa\_(n-1), aa\_(n-1)/@aa\_n/word:: wrd.*

[0068] The generated step queries are passed to composer 224, which accesses database 212 to retrieve the results for the step queries. In one embodiment, composer 224 may reduce the step queries to their canonical form and generate a hash key for the step query with methods as described above. Composer 224 references the hash key values in index 213 to  
25 retrieve the results from index 213. In one embodiment, the results may be one or more subtree IDs for the elements. In another embodiment, the results may be the elements corresponding to the step queries or all of the elements in each subtree.

[0069] Intersector 226 determines a result for the query using the step query results. The result is one or more subtrees that include all of the step queries. In one embodiment, one or  
30 more subtree IDs are returned. In one example, intersector 226 takes the intersection of the results of the step queries to produce a result that includes a result for query 219. An intersection of the results of the step queries may include some additional unwanted results.

A post-processing step may be performed by intersector 226 in which the unwanted results are eliminated. For example, the post-processing step matches each step query element of the intersection against the original query. After post-processing, the query result is outputted.

[0070] The optimization process implemented by optimizer 220 will now be described.

5 Optimizer 220 may optimize query 119 by rewriting query 119 in a form that may be used by step query generator 222 to generate optimized step queries. For example, optimizer 220 includes rewriting rules where queries 119 are rewritten in terms of the pre-computed step queries.

[0071] For example, optimizer 220 may rewrite path expressions of the form *aa\_1//aa\_2* by  
10 consulting a tree structure that represents the set of relations among the elements described for document 202 for which the query is intended. The *//* operator specifies the set of all nodes appearing below a given element node in the document tree. For example, *'A//B'* specifies the set of all the element nodes labeled 'B' which are strict descendants of 'A'. In one embodiment, optimizer 220 references elements described in XML schema where the  
15 schema includes, for each element that may occur in the document, a name, the type, the set of attributes, and the set of allowable constituent elements. The relations are represented in a graph, such as the one shown in Fig. 1, with one vertex for each element name, and one edge from an element to each possible constituent.

[0072] Given an XPath location path expression of the form *aa\_1//aa\_2*, optimizer 220  
20 attempts to determine a set of all possible sequences that interpolate the *"//"* (descendent-or-self::) step. The *'descendant-or-self::'* operator specifies the set of all nodes at or below a given node in the document tree. For example, *'A/descendant-or-self::B'* specifies the set of all descendants, including 'A', of the element node 'A'. For example, referring to Fig. 2, the expression *A//E* may be written as *(A/B/E union A/C/E)*. In some cases, the XPath location  
25 expression *A//E* may have an unlimited number of legal expansions. In this case, optimizer 220 does not attempt to rewrite the expression.

[0073] The optimized query is then passed to step query generator 222 for processing into step queries. The results to the step queries are retrieved as described above and the  
intersection taken by intersector 226. In post-processing, contiguous portions (maximal  
30 sequences of *"//"* separated steps) are processed as described above. Then, the results for the contiguous portions are then post-processed to verify the descendent relation by following parent links for the residual *"//"* steps. For example, the location path expression *A/C//B/E*

will be optimized as the pair of index queries A/C, B/E and for each node returned by the B/E parent links followed, parent links are followed to verify that some ancestor appears in the node set returned by A/C. The post-processing is done by creating an auxiliary hash index for the node ids occurring in A/C.

5 [0074] The post-processing step takes a sequence of subtree ids returned by the query composer and intersector, and scans these subtrees for the purpose of resolving general XPath location path expressions. A general XPath location path has the following syntax:

$$A[p_{11}]...[p_{1a}]/B[p_{21}]...[p_{2b}]/C[p_{31}]...[p_{3c}]/.../T[p_{k1}]...[p_{kt}]$$

or

10  $A[p_{11}]...[p_{1a}]/B[p_{21}]...[p_{2b}]/C[p_{31}]...[p_{3c}]/.../@T[p_{k1}]...[p_{kt}]$

[0075] Each of A, B, C, ..., T is an element label, and the last step T may be either an element label, an attribute label, or a non-element node type selector, e.g., 'text()' for text nodes, 'comment()' for comment nodes, and 'pi()' for processing instruction nodes. Each  $p_{ij}$  is a predicate expression that may involve as inputs additional (nested) XPath expressions, or  
 15 in the case of XQuery, variables from an enclosing scope. The post-processing step starts with a set of nodes corresponding to the last step. These are extracted from the subtrees appearing in the intersection of the pair step query results. This is called the 'candidate set'. For each node in the candidate set, the post-processing step moves backwards through the location path expression -- that is, the chain of ancestor nodes is obtained. (This may entail  
 20 additional access to the database subtree store.) For each ancestor chain, the location path expression is tested in its entirety, to verify that the element node labels correspond, and then that each sequence of step predicates  $p_{k1}, \dots, p_{kj}$  evaluates to 'true'. The post-processing algorithm is a 'generate-and-test' algorithm: for each element in the candidate set, a full path is generated by following parent links, and then this path is tested against the given location  
 25 path expression in its entirety. The post-processing generator does not evaluate predicates for candidates whose ancestor chain fails to match the node label pattern, (e.g.) A/B/C/.../T. For example, the location path query

$$A[p()] / B[q()] / C[r()] / @D$$

30 will be resolved as:

1. Composer 224 takes step queries and forms the set of index search queries:

Q1 : element-child-descendant-query("A", "B")

Q2 : element-child-descendant-query("B", "C")

Q3 : element-attribute-query("C", "D")

[0076] The first step query Q1 corresponds to the canonicalized term **A##B**, the second step  
 5 query Q2 corresponds to the canonicalized term **B##C**, and the third step query Q3  
 corresponds to the canonicalized term **C#@D**. Q1 and Q2 are 'descendant-queries', which  
 means that the queries specify a search among the descendants of nodes matching the given  
 pattern. In this example, the first step query 'xqe:element-child-descendant-query("A", "B")'  
 specifies a search among the descendants of **B** within subtrees containing the node pattern  
 10 '**A/B**' .

2. Intersector 226 receives a search of:

search(and-query(Q1, Q2, Q3), "C").

[0077] A search is performed for the intersection of Q1, Q2, and Q3, which returns a  
 15 sequence of nodes labeled **C**. The PostingLists for the canonicalized terms corresponding to  
 Q1, Q2, and Q3 are retrieved from index 213, and then scanned for common subtree ids. The  
 PostingList **skip list** structure is used to prune the search for common subtree ids. A  
 PostingList block will be skipped over in the event that the 'maxSubTreeID' stored in the  
 block is actually smaller than any of the currently smallest remaining subtree id in the other  
 20 PostingLists.

3. Post-processing:

[0078] For each node **c** labeled **C**, generate the ancestor path going back two steps, and check  
 if grandparent(**c**) = **A**, parent(**c**) = **B**. If not, discard **c**, and loop around to processes the next  
 node. If yes, then test the entire XPath expression

25 **A[p<sub>11</sub>]...[p<sub>1a</sub>]/B[p<sub>21</sub>]...[p<sub>2a</sub>]/C[p<sub>31</sub>]...[p<sub>3a</sub>]/@D**

by evaluating the predicate expressions from leftmost (highest) to rightmost (lowest) step, as  
 specified in the XPath standard. In this example, the step tests include a test for an attribute  
 node labeled '**D**' following '**C**'.

[0079] Fig. 10 illustrates a flow chart for a process for generating database 212 according  
 30 to one embodiment. In step S400, relationships among elements in a document are  
 computed. In step S402, step queries are generated from the relationships between the  
 elements. For example, one-step, two-step, three-step, and four-step queries are generated  
 from the relationship of elements.

[0080] In step S404, the step queries are reduced to their canonical form. In step S406, the process generates a hash key for each canonical form of the step queries. Additionally, in step S408, results for the step queries are generated from the relationship of elements. In one embodiment, the results represent one or more subtree IDs for the elements relating to each step query.

[0081] In step S410, the step query results and the corresponding hash keys are stored in database 212.

[0082] Fig. 11 is a flow chart of a process for generating a result for query 119 according to one embodiment. In step S500, query 119 is received by XQE 200 at query processor 218.

In step S502, a query may be optimized. In step S504, the optimized query is then reduced into step queries.

[0083] In step S506, a hash key is generated for each of the step queries. In step S508, database 212 is accessed and step query results are retrieved using the calculated hash key. For example, a PostingList may be retrieved.

[0084] In step S510, the intersection of the step query results is taken and a query result is generated from the intersection. For example, results from the PostingList are subtree IDs. The method determines matching subtree IDs where the matching subtree IDs would include the step queries associated with the step query results.

[0085] In S512, the query result may be post-processed. In step S514, the query result is outputted.

[0086] An example using an embodiment of the present invention will now be described. Fig. 12 depicts a PostingList 1000 that may be stored for the structure shown in Fig. 8 according to one embodiment of the present invention. A plurality of step queries 1002 are shown. Each step query represents a combination of nodes shown in Fig. 8 and one or more subtree IDs 1004 are associated with each step query 1002. For example, the step query "c/a" is found in the subtrees "10" and "30".

[0087] For discussion purposes, the relationships that cross subtrees are shown without any link nodes, for example,  $\langle e \rangle \rightarrow \text{link node}(c) \rightarrow \text{link node}(e) \rightarrow \langle c \rangle$  is represented as  $\langle e \rangle \rightarrow \langle c \rangle$ . Also, the upper node in the relationship is used to determine the subtree ID that is associated with step query. For example, the step query  $\langle e \rangle \rightarrow \text{link node}(c) \rightarrow \text{link node}(e) \rightarrow \langle c \rangle$  is associated with subtree "40".

[0088] Fig. 13 depicts PostingList 1000 with corresponding scores for each subtree ID according to one embodiment of the present invention. As shown, each step query 1002 and subtree ID 1004 pair has a score 1006 associated with it. Each score 1006 represents a numeric score that measures the relevance of step query 1002 to the step query in which it appears. In one embodiment, the score is computed by a function proportional to the number of occurrences of the term in the subtree divided by the total number of terms of any kind appearing in the subtree and may be normalized.

[0089] Figs. 14A - 14E depict PostingList structures for each subtree according to one embodiment of the present invention. Fig. 14A shows step queries 1100 that represent each step query found in the subtree represented by subtree ID "10". A frequency 1102 is shown for each step query 1100. A score 1104 is also shown for each step query 1100. For example, the step query "c/a" occurs twice in subtree 10 and has a score of .22.

[0090] Fig. 14B shows step queries 1100 that represent each step query found in the subtree represented by subtree ID "20". Fig. 14C shows step queries 1100 that represent each step query found in the subtree represented by subtree ID "30". Fig. 14D shows step queries 1100 that represent each step query found in the subtree represented by subtree ID "40". Fig. 14E shows step queries 1100 that represent each step query found in the subtree represented by subtree ID "50".

[0091] Using the above PostingLists described in Figs. 12-14, the following step query "b/c/a" may be queried. The query is broken down into the step queries of "b/c" and "c/a". Table I shows values that may be retrieved for PostingList 1000.

Step Query	Subtree ID → Score	Subtree ID → Score
b/c	20 → .08	50 → .18
c/a	10 → .22	30 → .18

Table I

The intersection of the subtree IDs for each step query is then taken. For example, the subtree IDs "20" and "50" are intersected with the subtree IDs "10" and "30". The intersection of these IDs is empty.

[0092] Although the intersection is empty, a further step may be taken to resolve the query. A query for step queries that may include link nodes (e.g., link node(c)) is then performed.

Thus, step query results for the step query "<b>→ link node (c)" may be retrieved. The results returned would be b/ link node (c) = subtree ID 20 → (subtree ID 10); subtree ID 50 → (subtree ID 20); and subtree ID 50 → (subtree ID 40). The above means that a "b/c" step query is linked across the subtrees 20/10, 50/20, and 50/40.

5 [0093] Table II represents the new results including the linked step query results.

Step Query	Subtree ID → Score	Subtree ID → Score	Subtree ID → Score
b/c	20 → .08	50 → .18	
c/a	10 → .22	30 → .18	
b/link node c	20(10)	50(20)	50(40)

Table II

[0094] The intersection of the three lists yields a set of candidate subtrees where the path b/c/a might occur, in this case, subtree ID 20 → subtree ID 10. As shown in Fig. 8, the path b/c/a is found in subtrees 20 and 10. The subtrees are then retrieved and examined to verify  
 10 the presence or absence of the path b/c/a. The path is then returned as the result of the query.

[0095] Although scores were not used in the above example, in one embodiment, scores may be used to determine the relevance of step query results. If many results are returned, the scores may be used to determine which step query results may be processed first. For example, if a subtree ID has a high score, then it may be more likely that the subtree  
 15 corresponding to the ID includes the elements of the query. Additionally, the subtree may be more relevant for the query. The scores may thus be used to prioritize processing and also to provide a list that represents the relevance of subtrees for a query.

[0096] In another example, system 200 may process results to determine if any false  
 20 positives are returned for the results. Using the query, "c/a/b", the following step queries are determined: "c/a" and "a/b". Table III depicts an example PostingList.

Step Query	Subtree ID → Score	Subtree ID → Score	Subtree ID → Score
c/a	10 → .22	30 → .18	
a/b	10 → .22	30 → .18	50 → .27

Table III

[0097] The intersection of the step query results for the two step queries yields the subtree IDs of "10" and "30". The subtree fragments corresponding to the subtree IDs are then retrieved. Each subtree fragment includes both step queries. However, the path in a fragment may not include the full query. For example, the "c/a" fragment should end with the "a/b" fragment. The subtree fragments are then traversed to determine if a fragment includes the query. Fig. 15A shows a false positive match and Fig. 15B shows a positive match. As shown in Fig. 15A, the "c/a" element does not connect directly to the "a/b" element. Thus, the fragment does not contain a "c/a/b" path. In Fig. 15B, a positive is shown as the path "c/a/b" is found in the fragment. This path may be returned as the result of the query.

10 [0098] In one embodiment, XQE 200 may be used to search for text in documents 202. A text search involves retrieving node sets (XML document fragments) that are relevant to a given set of terms. For example, a text search query may have the form: 'return all Citation nodes whose text content is relevant to the phrase "knee surgery".' Complex search query results may be reduced to the intersection of step query results in much the same way that location path queries may be reduced to the intersection of step queries followed by the generate-and-test post-processing step. A search for a set of terms { term\_1, term\_2, ..., term\_n } or a phrase "term\_1 term\_2 . . . term\_k" within the set of nodes with a given element A can be directly resolved by doing an index lookup for the terms: A#word(term\_1), A#word(term\_2), ..., A#word(term\_k), followed by an intersection of the results. For a phrase query, a post-processing step will verify that the terms appear contiguously in the subtrees.

[0099] XQE 200 indexes support full-text search across index 213. XQE 200 includes a set of built-in functions which resolve a variety of full-text queries, with methods:

- 25 (a) for constructing AND queries that specify a search within the intersection of any number of subsets specified by sub-queries,
- (b) for constructing OR queries that specify a search within the union of any number of subsets specified by sub-queries,
- (c) for constructing AND-NOT queries that specify a search within the set difference of two subsets specified by sub-queries,
- 30 (d) for constructing WORD queries that specify a search within elements whose text nodes contain a given phrase,
- (e) for constructing ELEMENT queries that specify a search within the set of

elements with a given QName,

(f) for constructing ELEMENT-VALUE queries that specify a search within the set of elements with a given QName whose full text value exactly matches a given phrase,

5 (g) for constructing ELEMENT-WORD queries that specify a search within the set of elements with a given QName whose text nodes contain a given phrase,

(h) for constructing ELEMENT-ATTRIBUTE queries that specify a search within the set of elements with a given QName which contain an attribute with a given QName,

10 (i) for constructing ELEMENT-ATTRIBUTE-VALUE queries that specify a search within the set of elements with a given QName which contain an attribute with a given QName, such that the attribute text exactly matches a given phrase,

(j) for constructing ELEMENT-ATTRIBUTE-WORD queries that specify a search within the set of elements with a given QName which contain an attribute with a given QName, such that the attribute text contains a given phrase,

15 (k) for constructing ELEMENT-CHILD queries that specify a search within the set of elements with a given QName that have a child with a given QName,

(l) for constructing ELEMENT-DESCENDANT queries that specify a search within the set of descendants of an element with a given QName,

20 (m) for constructing ROOT-ELEMENT-DESCENDANT queries that specify a search within the set of descendants of an element with a given QName whose parent is a document node with a given QName,

(n) for constructing ELEMENT-CHILD-DESCENDANT queries that specify a search within the set of descendants of an element with a given QName whose parent has a given QName,

25 (o) for constructing ELEMENT-ATTRIBUTE-DESCENDANT queries that specify a search within the set of descendants of an element with a given QName that has an attribute with a given QName,

(p) for constructing ELEMENT-ATTRIBUTE-VALUE-DESCENDANT queries that specify a search within the set of descendants of an element with a given QName that has an attribute with a given QName, such that the attribute text exactly matches a given phrase,

30 (q) for constructing URI queries that specify a search within the set of documents matching a given URI string.

A QName is a 'Qualified Name', which means a name of the form 'prefix:name', where **prefix** maps to some namespace URI, and **name** is any well-formed element or attribute name.

[0100] In one embodiment, complex text search queries are assembled by composition of the '*and-query*', '*or-query*' and '*and-not-query*' functions. The value of these functions is a  
 5 '*query value*', which represents a specification of a search pattern, which may be stored and evaluated at some subsequent point in the processing performed by XQE 200. The *query value* represents a delayed evaluation -- the *query value* specification determines a set of element sub-tree ids, but does not actually extract them from the database until passed to a function 'search'. The function 'search' may take two arguments: a *query value* and an  
 10 element QName, and evaluate the query specified by the *query value* argument returning a sequence of element sub-tree ids as specified by the QName argument. The QName argument may be an ancestor (or self) of the nodes returned by the *query value* specification.

[0101] In one embodiment, the sequence of nodes returned by the function search may be ordered by a 'relevance' score. The relevance of a node to the specified query is a complex  
 15 function that depends on the frequency the query terms appear in the text of the query nodes, the frequency the query terms appear across the entire database, and the quality score attached to a given node. The quality score is further described in Linblad IV-A.

[0102] In one embodiment, XQE 200 calculates the relevance of a node relative to any of the previously described precomputed text queries as the stored 'score' value in the  
 20 PostingList. Scores are composed through *and-query*'s and *or-query*'s by summation. The function *and-query* takes a sequence of any number of *query values* as an argument and returns a *query value* specifying a search matching all of the argument queries. The function *or-query* takes a sequence of any number of *query values* as an argument and returns a *query value* specifying a search matching any one of the argument queries. The function *and-not-*  
 25 *query* takes two *query value* arguments and returns a *query value* specifying a search matching the first but not the second argument query. Complex queries may be built by successive application of *and-query*, *or-query* and *and-not-query*. For example,

30 **and-query(or-query((element-word-query(QName("A")), "best"), element-word-query(QName("A")), "worst"))), element-word-query("A", "times"))),**

specifies a query for elements labeled "A", containing the term "times" and either one of the terms "best" or "worst". In addition, each of the *query value* functions can accept an argument specifying a relative weight for the query as a constituent of the composed query. For example,

5

```
and-query(or-query((element-word-query(QName("A")), "best", 0.7), element-word-  
query(QName("A")), "worst", 0.4))), element-word-query("A", "times", 0.9))
```

10 specifies a query for elements labeled "A", containing the term "times" and either one of the terms "best" or "worst", with the appearance of "worst" given relative weight 0.4, the appearance of "best" a relative weight of 0.7, and the appearance of "times" given a relative weight of 0.9. The relative weights are used when assigning an ordering to the result of a query.

15 [0103] Embodiments of the present invention provide methods for generating a pre-computed index that is used for generating a result for a query. Step queries are pre-computed and the results to these step queries generated and stored in the index along with the step queries. The step queries include a set of elements that are related in a parent-child relationship and may be used to generate a result for a query. Embodiments of the present invention receive a query and break the query into multiple step queries using elements from  
20 the path of the query. Results from these step queries are then retrieved from the index and the intersection of the retrieved results is taken to generate a result for the query. The result yields a location or elements that satisfy the query.

[0104] In one embodiment, relationships among elements in XML documents are computed and possible step queries that may be generated from the XML documents are  
25 computed. Because these step queries are pre-computed, XQueries that include the pre-computed step queries are satisfied in an efficient manner. Instead of traversing the hierarchical structure of the XML document on a node-by-node basis to find an element of the document, an index of pre-computed results for step-queries is used to generate a result for the query.

30 [0105] The above description is illustrative but not restrictive. Many variations of the invention will become apparent to those skilled in the art upon review of the disclosure. The

scope of the invention should, therefore, be determined not with reference to the above description, but instead should be determined with reference to the pending claims along with their full scope or equivalents.

WHAT IS CLAIMED IS:

- 1           1.       A method for generating a result for a query of a document of elements  
2 using pre-computed step queries and pre-computed step query results stored in a database, the  
3 method comprising:  
4           receiving the query, wherein the query comprises a path of elements in the  
5 document of elements;  
6           reducing the query into a plurality of step queries, wherein a step query  
7 comprises a relationship between a plurality of elements determined from a part of the path of  
8 elements;  
9           for each of the plurality of step queries, retrieving a pre-computed step query  
10 result for a step query in the plurality of step queries by querying the database using the step  
11 query, wherein the step query corresponds to a pre-computed step query for the pre-computed  
12 step query result; and  
13           generating the result for the query using the step query results.
- 1           2.       The method of claim 1, wherein generating the result comprises taking  
2 the intersection of the step query results.
- 1           3.       The method of claim 1, wherein the result of the query comprises a  
2 location in the document of elements that includes the path of elements for the query.
- 1           4.       The method of claim 1, wherein the result of the query comprises the  
2 path of elements for the query.
- 1           5.       The method of claim 1, further comprising optimizing the query,  
2 wherein optimizing the query comprises generating sequences from the path of elements that  
3 interpolate the path.
- 1           6.       The method of claim 1, wherein the plurality of step queries comprise  
2 at least one of a one-step query, two-step query, three-step query, and four-step query.
- 1           7.       The method of claim 1, wherein reducing the query into the plurality of  
2 step queries comprises reducing the query into at least one two-step query.
- 1           8.       The method of claim 1, wherein reducing the query into the plurality of  
2 step queries comprises reducing the query into at least one three-step query.

1           9.     The method of claim 1, further comprising  
2           computing a hash key for queries in the pre-computed step queries and  
3     plurality of step queries; and  
4           storing the hash keys for the pre-computed step queries and the corresponding  
5     pre-computed step query results in the database.

1           10.    The method of claim 9, wherein retrieving the pre-computed step  
2     query result comprises using the stored hash keys for the step queries to retrieve the pre-  
3     computed step query results corresponding to the hash keys.

1           11.    The method of claim 9, wherein the step query results comprise a ID  
2     for one or more elements in the document of elements.

1           12.    The method of claim 9, further comprising post-processing the  
2     intersection of the step query results to generate the result for the query.

1           13.    The method of claim 12, wherein post-processing the result comprises  
2     matching each step query in the step query results to the query.

1           14.    The method of claim 9, wherein the relationship between the plurality  
2     of elements comprises a parent/child relationship.

1           15.    The method of claim 9, wherein the document of elements comprise an  
2     XML document.

1           16.    The method of claim 9, wherein elements in the document of elements  
2     comprise at least one of element, word, attribute, and string elements.

1           17.    A method for creating a database of step queries and step query results  
2     for a document of elements, the method comprising:  
3           determining relationships between a plurality of elements from the document  
4     of elements;  
5           generating step queries from the relationships;  
6           generating step query results for the step queries, wherein a step query result  
7     for a step query corresponds to one or more elements in the document of elements for the step  
8     query; and

9 storing the step queries and corresponding step query results in the database,  
10 wherein the stored step query results are usable to generate a result for a main query, wherein  
11 the main query can be reduced to a plurality of step queries that correspond to the stored step  
12 queries.

1 18. The method of claim 17, further comprising generating an index for the  
2 step queries, the index pointing to the corresponding step query results for each step query.

1 19. The method of claim 17, wherein the step query results comprise a ID  
2 for one or more elements in the document of elements.

1 20. The method of claim 17, wherein the plurality of step queries and  
2 corresponding step query results are stored in a PostingList.

1 21. The method of claim 17, wherein the step queries comprise at least one  
2 of a one step query, two step query, three step query, and four step query.

1 22. The method of claim 17, wherein the document of elements comprise  
2 an XML document.

1 23. The method of claim 17, wherein elements in the document of  
2 elements comprise at least one of element, word, attribute, and string elements.

1 24. The method of claim 17, wherein generating step queries comprises  
2 reducing a step query to a canonical form.

1 25. The method of claim 17, wherein the relationship between the plurality  
2 of elements comprises a parent/child relationship.

1 26. The method of claim 17, wherein storing the step queries comprises:  
2 generating a hash key for every step query; and  
3 storing the hash key for the step queries in the database.

1 27. A query processor for processing a query for a document of elements,  
2 the processor comprising:  
3 a document processor configured to receive the document of elements and pre-  
4 compute a plurality of step queries and corresponding step query results from the document  
5 of elements;

6 a database for storing the pre-computed plurality of step queries and  
7 corresponding step query results; and  
8 a query processor configured to receive the query, generate a plurality of step  
9 queries from the query, and generate a result for the query using the step query results  
10 retrieved from the database that correspond to the plurality of step queries.

1 28. The query processor of claim 27, wherein the document processor  
2 comprises a step query generator configured to generate the pre-computed plurality of step  
3 queries.

1 29. The query processor of claim 27, wherein the document processor  
2 comprises a canonicalizer configured to generate a canonical form of the pre-computed  
3 plurality of step queries.

1 30. The query processor of claim 27, wherein the document processor  
2 comprises a hash key generator configured to generate a hash key for each of the pre-  
3 computed plurality of step queries.

1 31. The query processor of claim 27, wherein the document processor  
2 comprises a step query result generator configured to generate step query results for the pre-  
3 computed plurality of step queries.

1 32. The query processor of claim 27, wherein the step query results  
2 comprise one or more identifiers corresponding to one or more elements in the document of  
3 elements.

1 33. The query processor of claim 27, wherein the query processor  
2 comprises an optimizer configured to optimize the query.

1 34. The query processor of claim 27, wherein the query processor  
2 comprises a step query generator configured to generate a plurality of step queries from the  
3 query.

1 35. The query processor of claim 27, wherein the query processor  
2 comprises a composer to retrieve the step query results retrieved from the database that  
3 correspond to the plurality of step queries.

1           36.     The query processor of claim 27, wherein the query processor  
2 comprises an intersector configured to take the intersection of step query results retrieved  
3 from the database that correspond to the plurality of step queries.

1           37.     The query process of claim 27, wherein the document of elements  
2 comprises an XML document.

1           38.     A method for processing queries for a document of elements, the  
2 document including a plurality of subsections, each subsection including at least a portion of  
3 elements in the document, the method comprising:

4                 receiving a query for a path of elements in the document of elements;

5                 determining a plurality of step queries from the query, each step query  
6 including at least a part of the path of elements;

7                 for each step query in the plurality of step queries, determining one or more  
8 subsections that include elements that correspond to a step query; and

9                 determining at least one subsection that includes the path of elements of the  
10 query.

1           39.     The method of claim 38, further comprising generating a result for the  
2 query using the at least one subsection.

1           40.     The method of claim 39, wherein the result comprises a location where  
2 the path of elements is stored.

1           41.     The method of claim 39, wherein the result comprises the path of  
2 elements.

1           42.     The method of claim 38, wherein determining one or more subsections  
2 comprises determining a subsection identifier for each of the one or more subsection.

1           43.     The method of claim 38, further comprising:

2                 determining a relevance score for each of the one or more determined  
3 subsections; and

4                 using the relevance scores for the one or more determined subsections to  
5 determine the at least one subsection that includes the path of elements.

1                   44.     The method of claim 38, further comprising verifying that the at least  
2 one subsection that includes the path of elements actually includes the path of elements using  
3 the query.

1                   45.     The method of claim 38, further comprising:  
2                   determining, for each of the one or more determined subsections, how many  
3 times an instance of a step query appears in a subsection ; and  
4                   using the frequency to determine the at least one subsection that includes the  
5 path of elements.



1/13

30

```
<citation>
  <title>Cerisent XQE</title>
  <author>
    <last>Pedersen</last>
    <first>Paul</first>
  </author>
  <abstract> The document describes an XML
               search and query system
  </abstract>
</citation>
```

FIG. 1

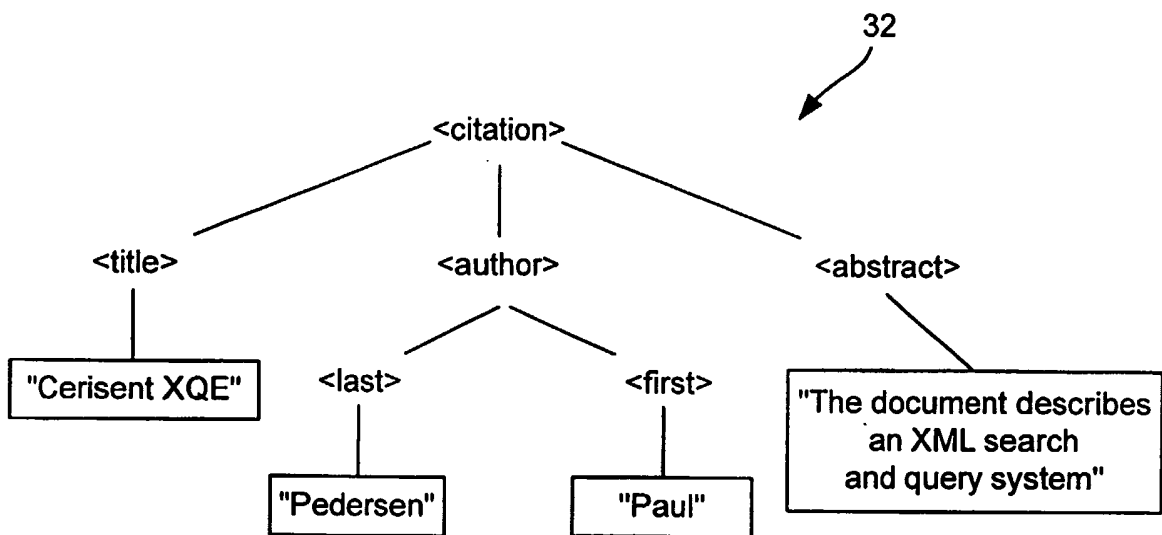


FIG. 2A




FIG. 2B




FIG. 3

< b K = "v" > node text</b>

FIG. 4A




FIG. 4B




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9

Query Result

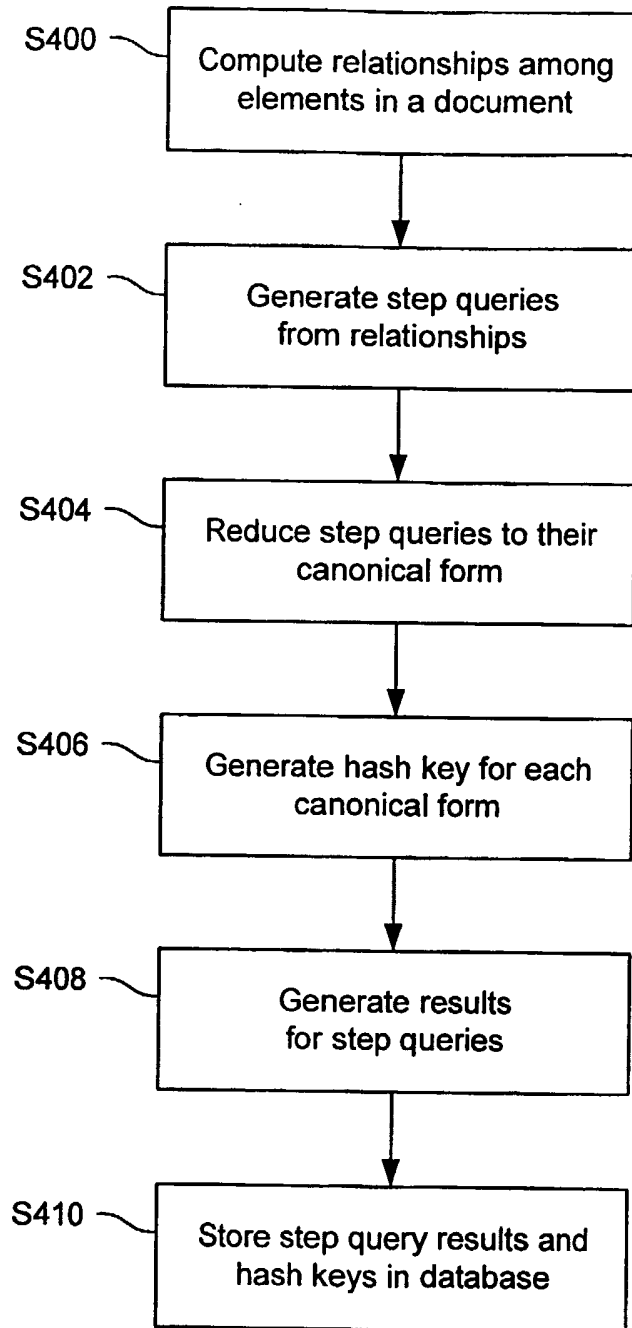


FIG. 10

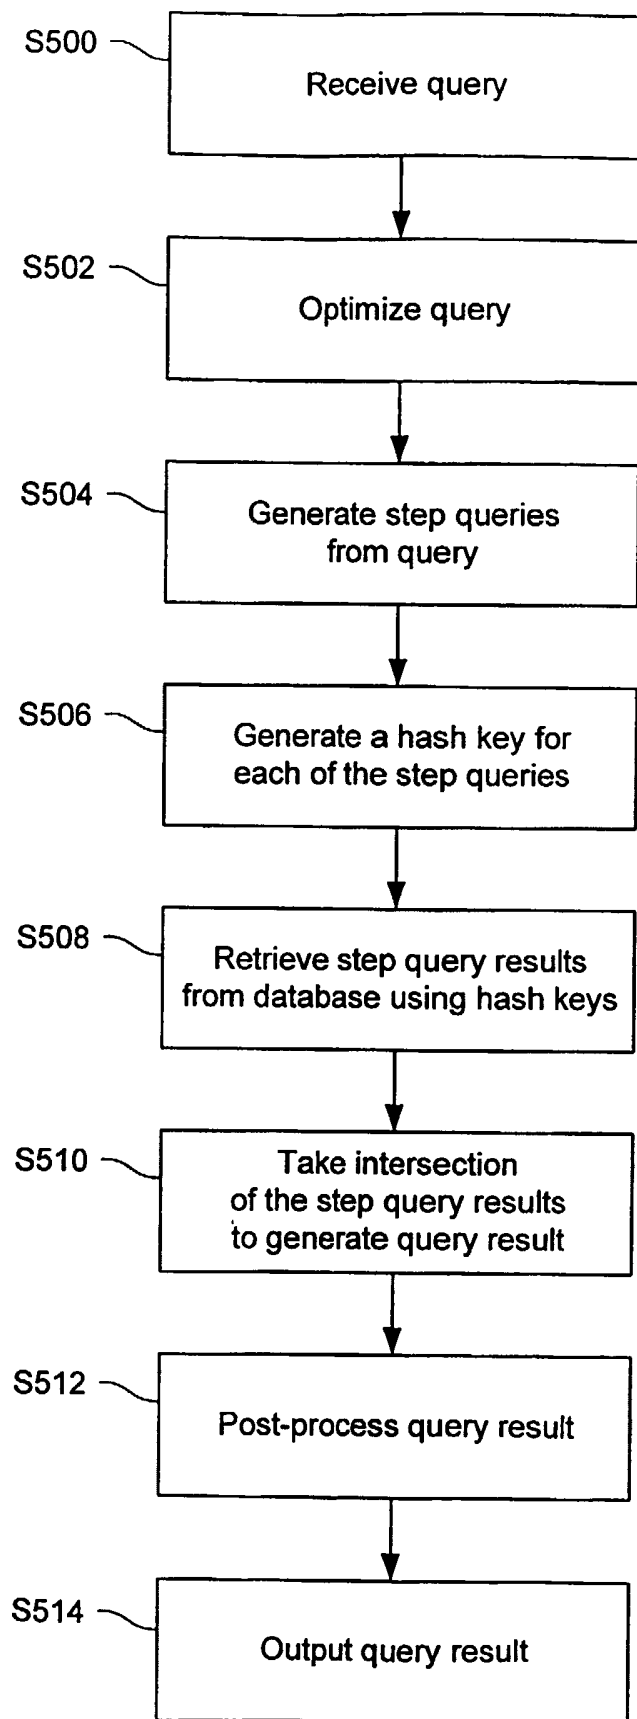


FIG. 11

1000  
↙

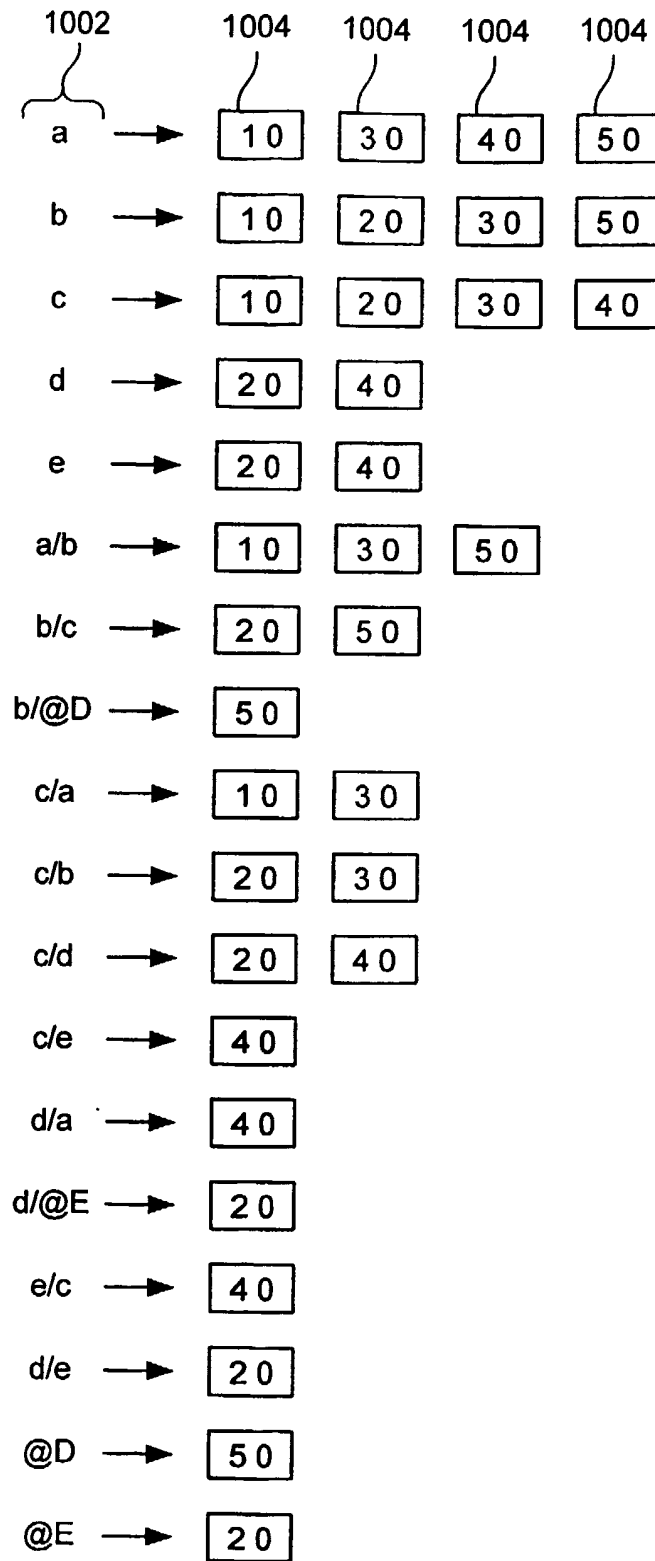



FIG. 12



		1004	1006	1004	1006	1004	1006	1004	1006
1002	a	10	.22	30	.18	40	.13	50	.09
b	10	.22	20	.08	30	.27	50	.27	
c	10	.11	20	.08	30	.09	40	.13	
d	20	.08	40	.13					
e	20	.17	40	.25					
a/b	10	.22	30	.18	50	.27			
b/c	20	.08	50	.18					
b/@D	50	.09							
c/a	10	.22	30	.18					
c/b	20	.09	30	.09					
c/d	20	.08	40	.13					
c/e	40	.13							
d/a	40	.13							
d/@E	20	.08							
e/c	40	.13							
d/e	20	.17							
@D	50	.09							
@E	20	.08							

FIG. 13



11/13

Step Queries 1100	Frequency	Score
a	2	.22
b	2	.22
c	1	.11
c/a	2	.22
a/b	2	.22

FIG. 14A

Step Queries 1100	Frequency	Score
b	1	.08
c	1	.08
d	1	.08
e	2	.17
@E	1	.08
c/b	1	.08
c/d	1	.08
d/e	2	.17
d/@E	1	.08
b/c	1	.08

FIG. 14B

Step Queries 1100	Frequency	Score
a	2	.18
b	3	.27
c	1	.09
a/b	2	.18
c/a	2	.18
c/b	1	.09

12/13


FIG. 14C

Step Queries 1100	Frequency	Score
c	1	.13
d	1	.13
e	2	.25
c/d	1	.13
c/e	1	.13
d/e	1	.13
e/c	1	.13

FIG. 14D


Step Queries 1100	Frequency	Score
a	1	.09
b	3	.27
@D	1	.09
a/b	3	.27
b/c	2	.18
b/@D	1	.09

FIG. 14E



False Positive

FIG. 15A



Positive

FIG. 15B



**INTERNATIONAL SEARCH REPORT**

International application No.  
PCT/US03/18873

<b>A. CLASSIFICATION OF SUBJECT MATTER</b>		
IPC(7) : G06F 17/30 US CL : 707/3, 102		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols) U.S. : 707/3, 102		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) NPL, STIC, IEEE		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y, P	US 6,510,406 B1 (MARCHISIO) 21 January 2003 (21.01.2003), figure 1, column 5, lines 8-57.	1-45
Y, P	US 6,438,540 B2 (NASR et al) 20 August 2002 (20.08.2002), figure 1, especially item 150, column 2, lines 10-12, 20-22 and 25-27.	1-45
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.		
* Special categories of cited documents:		
"A"	document defining the general state of the art which is not considered to be of particular relevance	"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
"E"	earlier application or patent published on or after the international filing date	"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
"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)	"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
"O"	document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed	
Date of the actual completion of the international search	Date of mailing of the international search report	
14 October 2003 (14.10.2003)	29 OCT 2003	
Name and mailing address of the ISA/US	Authorized officer	
Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450 Alexandria, Virginia 22313-1450	Safet Metjahic <i>Peggy Harwood</i>	
Facsimile No. (703)305-3230	Telephone No. 703-305-3900	