

## AN EFFICIENT AND EXTEND B-TWIG PATTERN FOR XML QUERY PROCESSING

Santhosh Kumar KP\*

T.Kumesh\*\*

### Abstract—

Twig pattern matching is a critical operation for XML query processing, and the holistic computing approach has shown superior performance over other methods. Since Bruno et al. introduced the first holistic twig join algorithm, TwigStack, numerous so-called holistic twig join algorithms have been proposed. Yet practical XML queries often require support for more general twig patterns, such as the ones that allow arbitrary occurrences of an arbitrary number of logical connectives (AND, OR, and NOT); such types of twigs are referred to as B-twigs (i.e., Boolean-Twigs) or AND/OR/NOT-twigs. We have seen interesting work on generalizing the holistic twig join approach to AND/OR-twigs and AND/NOT-twigs, but have not seen any further effort addressing the problem of AND/OR/NOT-Twigs along with XOR twig at the full scale, which therefore forms the main theme of this paper. In this paper, we investigate novel mechanisms for efficient B-twig pattern matching. In particular, we introduce “B-twig normalization” as an important first-step in our approach toward eventually conquering the complexity of B-twigs, and then present BTwigMerge the first holistic twig join algorithm designed for B- twigs. Both analytical and experimental results show that BTwigMerge is optimal for B-twig patterns with AD (Ancestor-Descendant) edges and/or PC (Parent-Child) edges.

**Index Terms:** XOR twig, Query processing, database management, XML data querying, twig join, Boolean twig, logical predicate

\* PG Student, Department of Computer Science and Engineering, PSN Engineering College, Tirunelveli-627152, India

\*\* Assistant Professor of the Department, Department of Computer Science and Engineering, PSN Engineering College, Tirunelveli-627152, India

## 1. INTRODUCTION

AN XML database stores a collection of data trees. An XML query describes a tree-shaped search pattern, which is often referred to as a twig pattern [5], with additional query conditions (if any) described as predicates on the three nodes. XML queries thus are called tree queries or twig queries. Answering a twig query is essentially to find all matching instances from a database that match the twig pattern implied by the query and satisfy all additional predicates (if any) in the query. A naive way of finding the matches for a twig pattern is to scan the database (usually for many times). A better way uses structural joins [14], [4] in a bulk way to compute the matches for each individual edge, and then “stitch” the matches found for individual edges together to form the answers for the whole twig. This approach typically creates large sets of unused intermediate results, even when the final result set is pretty small. Yet, a much more efficient approach, called holistic twig join, computes the matches for the whole twig in a holistic way so that irrelevant intermediate results (which need be output and input, and thus are most detrimental to query performance) can be avoided. The first holistic twig join algorithm, TwigStack, was proposed by Bruno et al. [5] in 2002. Since then the “holistic join” approach has been done broadly extended by numerous followers [6], [8], [7], [9], selects the authors who have papers either titled “TwigJoin” or published in SIGMOD 2006. This query contains both OR and AND operations. The next query, finds papers that do not have references. This query contains a NOT operation. A twig that may contain arbitrary combination of ANDs, ORs, and NOTs, is referred to as an AND/OR/NOT-twig or Boolean twig (or simply B-twig).

The importance of B-twigs for XML queries is obvious and well recognized [7], [13]. So far, we only see that Jiang et al. [7] studied the holistic twig join issue for AND/OR-twigs (i.e., twigs with only AND and OR predicates) and Yu et al. [13] tackled the problem for AND/NOT-twigs (i.e., twigs with only AND and NOT predicates). There is no integral method ever reported facing the full challenge of holistic B-twig computing. The challenge with full B-twigs lies in the arbitrary occurrences of an arbitrary number of AND/OR/NOT predicates in a B-twig (we refer to this challenge as the “double arbitrariness” challenge of B-twigs). This challenge makes programmatic handling of B-twigs in the framework of holistic computing extremely hard (if not impossible). We have made numerous years of effort on conquering the full challenge of B-twigs. We could not easily sort out the complication caused by the “double

arbitrariness” of B-twig, and thus could not systematically and programmatically solve the problem of B-twig with a nice algorithm. However, our effort has helped us gain in-depth insight into the challenge of holistic B-twig pattern computing. We have made numerous years of effort on conquering the full challenge of B-twig. We could not easily sort out the complication caused by the “double arbitrariness” of B-twig, and thus could not systematically and programmatically solve the problem of B-twig with a nice algorithm. The severity of this challenge, we believe, there are so many ways holistic join solution for B-twig that has not been developed nearly 9 years after Bruno et al. first proposed the promising holistic join approach [5] that afterward quickly inspired the solutions for AND/OR-twig [7] and AND/NOT-twig [13] separately proposed by different researchers. From AND/OR-twig and AND/NOT-twig to full B-twig appears to be just one step, however, as the complexity implied by the double arbitrariness blows up, a holistic join approach for full B-twig cannot be simply obtained from combining the methods separately designed for AND/OR-twig and AND/NOT-twig. Rather, a more creative strategy with more powerful supporting mechanisms must be invented for B-twig. Solving the challenge of holistic B-twig computing has both practical and academic significance. From the practical perspective, this effort helps to mature the promising holistic twig join approach and can immediately find use in real XML query applications; from the academic side, it solves an important technical problem and the obtained result can be generalized to any data sources incarnating a tree data model (while XML is just one use case of the general tree data model).

We are thus motivated to sort out the complication involved in holistic computing of B-twig pattern matches. In this paper, we present our complete approach, including the techniques we developed for systematically solving holistic B-twig computing. The contributions of our work reported here can be summarized as follows:

- We propose a novel facility, i.e., B-twig normalization, which serves as the first milestone in our approach toward eventually solving the increased complexity of B-twig.
- We expound a sound method for automatically performing B-twig normalization, which is an important prestep in our overall approach.
- We present BTwigMerge, the first holistic join algorithm ever designed for (normalized) B-twig, including numerous original supporting mechanisms.

- BTwigMerge performs optimal matching [5] for both AD (Ancestor-Descendent) edges and PC (Parent-Child) edges, while prior algorithms claim optimality only for AD edges.

The remainder of this paper is organized as follows: Section 2 reviews related work. Section 3 sets forth the preliminaries for the subsequent discussion, including data model, B-twig representation, and normalization. Section 4 presents our algorithm, BTwigMerge, including its various supporting functions (each implements an important supporting mechanism). Section 5 provides experimental results, demonstrating the superiority of our approach and algorithm.

## 2 RELATED WORK

Twig pattern matching is a core operation in the XML query processing. Naive navigation (or pointer-chasing), structural joins, and holistic twig joins have all been studied for twig pattern matching. In the following, we review representative works on structural joins and particularly on holistic twig joins.

The first structural join (called containment join) algorithm was proposed by Zhang et al. [14], which extends the traditional merge join to multipredicate merge join (MPMGJN). Al-Khalifa et al. [4] later proposed two families of structural join algorithms, i.e., tree-merge and stack-based structural joins, as primitives for XML twig query processing. In 2002, Bruno et al. [5] first proposed the holistic twig join approach for XML twig queries in order to overcome the drawback of structural joins that usually generate large set of unused intermediate results. Bruno et al. designed the first holistic twig join algorithm, named TwigStack, which is optimal for twigs with only AD edges (but not with PC edges). The work of Lu et al. [9] aimed at making up this flaw and they presented a new holistic twig join algorithm, TwigStackList, in which a list structure is used to cache limited elements in order to identify a larger optimal query class. Chen et al. [6] studied the relationship between different data partition strategies and the optimal query classes for holistic twig joins. Lu et al. [10] proposed a new labeling scheme, called extended Dewey, and an interesting algorithm, named *TJFast*, for efficient processing of XML twig patterns. Unlike all previous algorithms based on region encoding, to answer a twig query, *TJFast* only needs to access the labels of the leaf query nodes. The result of Lu et al. [10] includes enhanced functionality (can process limited wildcard), reduced disk access, and increased total query performance. The same group [11] also studied efficient processing for ordered XML twig patterns using their region encoding scheme.

In an ordinary twig, the multiple sibling nodes under a common parent node automatically signify the AND logic relationship among them, and all previously proposed holistic twig join algorithms already support this implied AND logic in their implementation schemes. Users would take all the three commonly used logical predicates, AND, OR, and NOT, as granted facilities in formulating their XML queries and thus would expect full support from a query engine for unlimited use of all these predicates in their XML queries. Jiang et al. [7] made the first effort toward incorporating support for OR predicates into the holistic twig join approach pioneered by Bruno et al. [5], and Yu et al. [13] made effort for supporting NOT predicates in XML twig queries. Jiang et al. [7] presented an interesting framework for holistic processing of AND/OR-twigs based on the concept of OR-block. With resort to OR-blocks, an AND/OR-twig is transformed to an AND-only twig carrying special OR-blocks. This work is inspiring to us—we find it is possible to substantially extend their framework so that the NOT logic can be seamlessly incorporated. Nevertheless, this work is not straightforward, but requires creative reinvention of the “wheels.” In order to harness the complexity of B-twigs, we resort to B-twig normalization; then based on normalized B-twigs, we are able to extend and adapt the OR-block concept with new supporting mechanisms for handling the NOT predicates involved in B-twigs.

The recent publication of Xu et al. [12] proposed another interesting algorithm that claims to be able to efficiently compute the answers to XML queries without holistically computing the twig patterns—the answers obtained contain individual elements corresponding to designated output query nodes. So basically this work does not belong to the category of holistic twig join algorithms. But what is interesting of their work [12] is the proposed path-partitioned element encoding scheme, which bears efficiency potential and may be considered in the future for further improving the performance of holistic B-twig pattern matching.

### 3 PRELIMINARIES

In this section, we first address the data model issue, B-twig representation and normalization, and then introduce the notations and operations needed in our subsequent discussion.

### 3.1 Data Model

We adopt the general perspective [5] that an XML database is a forest of rooted, ordered, and labeled trees, each node corresponds to a data element/value, and each edge represents an element-sub element or element-value relation. The order among sibling nodes implicitly defines a total order on the tree nodes. Node labels are important for efficient processing of a twig pattern as properly designed node labels may leave out the necessity of accessing the node contents during query evaluation. This is especially true with twig pattern matching, which is at the core of XML query processing. Node labels typically encode the region information of data elements that reflects the relative positional relationships among the elements in the source data file. We assume a simple encoding scheme using a triplet region code—(start; end; level)—which is assigned to each data element in a tree database as a label. When multiple documents are present, the document-id is added to the labels to differentiate the documents. Region code can be conveniently obtained through preorder document-tree traversing.

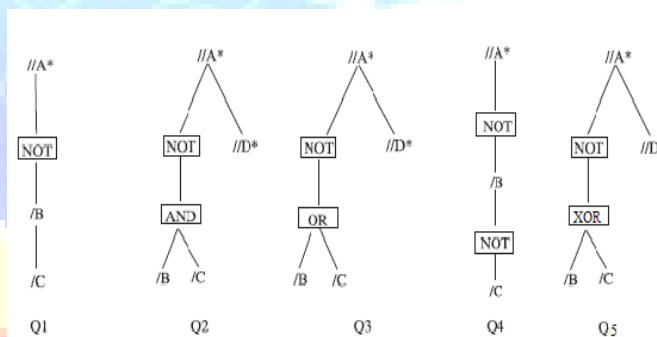
### 3.2 Tree Representation

Each XML query implies a twig pattern, small, or large. The smallest twig may contain just a single node, but a typical twig usually comprises a number of nodes. The target of our investigation is the B-twigs that allow arbitrary combination of AND, OR, and NOT predicates, of which each may have multiple occurrences. Each B-twig may consist of two general categories of nodes: ordinary query nodes standing for element types (or tags) and special connective nodes denoting logical predicates—AND, OR, and NOT. More specifically, we represent a B-twig using the following specific types of nodes:

- *QNode*. An ordinary query node, associates to an element type (or tag name) in a tree database. For programmatic purposes (as in [7]), a *QNode* records its location step axis “//” or “/” for edge test, and a tag name for node test. Therefore, the content of a *QNode* takes the general format of “/tag” or “//tag.” A non-root *QNode* in a B-twig may be conveniently called a d-child or c-child (of its parent) depending on whether a “//” or a “/” symbol is recorded in the *QNode*’s content (Notice that in the sequel we may not always show the location step axes in our illustrations when the emphasis is on something else).

- *ANode*. An AND predicate node, always takes the text “AND” as its content. It connects two or more child sub trees through the AND logic.
- *ONode*. An OR predicate node, always takes the text “OR” as its content. It connects two or more child sub trees through the OR logic.
- *NNode*. The NOT predicate node, always takes the text “NOT” as its content. Functionally, a *NNode* negates the predicate denoted by the sub tree immediately underneath it. A *NNode* is commonly combined with the node underneath it in the B-twig, forming a composite node. We have the following three kinds of composite nodes related to NOT
- *ZNode*. An XOR predicate node, always takes the text “XOR” as its content. It connects two or more child sub trees through the XOR logic.

Fig. 1. Example twigs involving NOT



-*NQNode*: the combined form of a NOT node with a subsequent *QNode* child (such combination just causes the representation of a B-twig more compact, and does not affect the semantics or interpretation of the twig pattern). For example, in the query Q1 (shown in Fig. 1), the NOT and the subsequent child *QNode* “/B” can be combined and replaced by a single *NQNode* with content “□/B.”

-*NANode*. The combined form of a NOT node with its sole *ANode* child.

-*NONode*. The combined form of a NOT node with its sole *ONode* child.

-*NZNode*. The combined form of a NOT node with its sole *ZNode* child.

We could also have the fourth type of composite node that represents a NOT node combined with another (child) NOT node (i.e., a double negation node that could be named *NNNode*). As the net effect of double negations is the same as no negation at all, double negation nodes are not actually used in our representation for B-twigs. From now on, we generally refer to *QNodes* and *NQNodes* as query nodes, and other (plain or composite connectives) nodes in a B-twig as no query nodes.

With the above mechanisms introduced, our representation scheme for B-twigs is apparently a superset of what can be represented by the scheme adopted by Jiang et al. [7] for the simpler AND/OR-twigs. Considering NOT as a new element added to B-twigs (comparing to the AND/OR-twigs studied in [7]), we next illustrate the four typical cases that the NOT predicates may appear in an XML twig query. The following four queries exemplify these four representative cases:

Q1. A[NOT/B/C]

Q2. A[NOT (/B AND/C)]//D

Q3. A[NOT (/B OR/C)]//D

Q4. A[NOT/B[NOT/C]]

Q5. A[NOT[B XOR/C]]

### 3.3 Edge test Algorithm

The main structure of function *edgeTest* (and function *nEdgeTest* as well) is a while loop, which at the first glimpse appears unnecessary, but (at lines 7 and 8, see Fig.2) brings an important optimization—fast skipping noncontributing elements in stream  $T_q$  until the cursor moves over the range of the parent element  $e$ . (This “fast skipping” has the effect of instant performance optimization, otherwise those noncontributing elements will stay in their streams causing extra iterations and consuming extra CPU time.). The location step axis is obtained from the content of the child query node. The implementation of function *nEdgeTest* relies on repeated calls to function *edgeTest* (see Fig.3). The implementation of function *ONodeTest* (see Fig. 4) almost straightforwardly follows Definition. It is based on *edgeTest*, *nEdgeTest*, and yet another function, *hasExtension*, that realizes.

#### FUNCTION *edgeTest*( $e, q$ )



```

1: while not  $end(Cq)$  do
2:   if  $e.start < Cq \rightarrow start$  and  $e.end > Cq \rightarrow end$ 
then
3:   if  $q.axis == '/'$  then
4:     return true
5:   else if  $e.level == Cq \rightarrow level - 1$  then
6:     return true
7:   if  $Cq \rightarrow end < e.end$  then
8:      $Cq \rightarrow advance()$ 
9:   else
10:    break
11: end while
12: return false

```

Fig.2. Function *edgeTest*.

Holistic twig joins typically disallow backtracking of stream cursors to guarantee linear time complexity. Notice that the evaluation of a NOT predicate involved in a B-twig requires disproving all elements in the negated stream (associated to the query node negated by the NOT predicate). This seems to imply that we need to scan to the very end of the negated stream to disprove all elements, and then backtrack the stream cursor to get ready for reevaluating the subsequent parent elements. In fact, such backtracking can be avoided.

#### FUNCTION *nEdgeTest*(*e,q*)

```

1: while not  $end(Cq)$  do
2:   if  $edgeTest(e,q) == true$  then
3:     return false
4:   else if  $Cq \rightarrow end < e.end$  then
5:      $Cq \rightarrow advance()$ 
6:   else
7:     break
8: end while
9: return false

```

Fig.3. Function *nEdgeTest*

Let's exemplify this: assuming we are processing the B-twig sub expression  $X \text{ NOT } Y$ , and elements  $x_i$  and  $y_j$  are currently under the cursor of their respective streams,  $T_X$  and  $T_Y$ , there are two cases to consider now: 1)  $y_j$  happens to fall within the region of  $x_i$ , this immediately disproves  $x_i$  and the evaluation immediately returns; 2)  $y_j$  is not in the range of  $x_i$ , this leads to two possible subcases: a)  $y_j$  is ahead of  $x_i$ —then we just advance cursor  $C_Y$  and start the next iteration to evaluate with the next element following  $y_j$ ; b)  $y_j$  falls behind  $x_i$ , this is enough to qualify  $x_i$ , since all subsequent elements after  $y_j$  (if any) can only be farther away from the coverage of  $x_i$  due to the sort order of the elements in stream  $T_Y$ . In all cases, advancing of the stream cursor  $T_Y$  never goes beyond the range covered by  $x_i$ , and backtracking is never needed. The code in Fig.4 embodies the idea discussed above.

**FUNCTION *ONodeTest*(e,n)**

```

1: for each  $n_i$  in  $P(n)$  do
2:   if isLeaf( $n_i$ ) and isQNode( $n_i$ )
3:     replace  $n_i$  by edgeTest( $e, n_i$ )
4:   else if isLeaf( $n_i$ ) and isNQNode( $n_i$ )
5:     replace
 $n_i$  by (edgeTest( $e, n_i$ ) and hasExtension( $n_i$ ))
6:   else /*  $n_i$  is a non leaf QNode
7:     replace  $n_i$  by edgeTest( $e, n_i$ )
8:   end for
9: evaluate  $P(n)$  and return the result

```

Fig.4. Function *ONodeTest***4 AHOLISTIC B-TWIG JOIN ALGORITHM**

With the supporting mechanisms set forth in the preceding sections, we now present our novel holistic B-twig join algorithm, *BTwigMerge*, in this section.

#### 4.1 *BTwigMerge*: The Main Algorithm

The structure of our main algorithm, *BTwigMerge*, as shown in Fig.5, is not much different from most other holistic twig join algorithms [5], [6], [8], [7], [9], [10], [11],[13]. It is a merge-based, two-phase algorithm. However, as we confront a different, more complex problem of B-twig that was not considered by previous holistic algorithms, the processing strategy of our *BTwigMerge* must be accordingly different. The difference mainly lies in the key supporting function, *GetQNode* (detailed in the next Section). In addition to feeding the main algorithm the next query node to be processed, our *GetQNode* function thoroughly investigates the candidacy of the elements in the input streams and guarantees that for the next *QNode* returned to the main algorithm the current element in the associated stream is fully qualified, i.e., the element satisfies all relevant criteria (including all predicates and edge tests). Functionally, algorithm *GTwigMerge* [7] is the closest to our *BTwigMerge*, but at the main algorithm level, *BTwigMerge* is more concise: with each valid query node *q* returned by *GetQNode* (the validity is checked at line 3, see Fig.5), *BTwigMerge* cleans up relevant stacks (lines 6 and 7), moves the element associated to *q* from stream to stack if it is not an output leaf (at lines 8 to 10), otherwise (*q* is an output leaf) outputs the path solutions currently on the stacks (lines 9 and 10). It is worth to point out that *BTwigMerge* does not explicitly process OR and any other predicates at the main algorithm level (differing from *GTwigMerge* [7]). Instead, all critical processing logics are encapsulated in the key supporting function, *GetQNode*, and other lower level supporting functions. *GetQNode* also checks whether each involved PC or AD edge is, respectively, satisfied by the stream head element associated to *q* that is to be returned to the main algorithm as the next *QNode* for processing. Therefore, our *BTwigMerge* achieves matching optimality not only with AD edges but also with PC edges (this is in strong contrast to all previous holistic algorithms including *TwigStack* [5] and *GTwigMerge* [7], etc.).

FUNCTION ORBlockMax(*n*)

1:  $q_0 = 0$  /\* refers to *QNode* \*/

2: if *isNQNode*(*n*) then

```

3 return 0
4: else if isQNode(n) and isLeaf(n) then
5: return n
6: else
7:   if isQNode(n) then
8:      $q_0 = n$ 
9:     for each  $n_i \in \text{children}(n)$  do
10:       $q_i = \text{ORBlockMax}(n_i)$ 
11:    end for
12:   if isONode(n) then
13:     return  $\text{argmin}_{q_i}\{e_i, \text{start}\}$ 
14:   else
15:     return  $\text{argmax}_{q_i}\{e_i, \text{start}\}$ 

```

Fig.6. Function *ORBlockMax*.

Some important features of *BTwigMerge* are highlighted as follows:

1. *BTwigMerge* receives (from *GetQNode*) either a valid output *QNode* *q* or an invalid *QNode*, denoted by null. An invalid *QNode* is typically generated by *GetQNode* when a non-top level recursive call into this function fails to find a *QNode* associated with a fully qualified element. But since non-contributing elements encountered during this process have been skipped, the main algorithm quickly jumps to its next iteration (at line 4) to start a new call to *GetQNode* for getting the next valid *QNode*.
2. No stacks are allocated for no output *QNodes*, nor for any output leaves (*QNodes*) since the contributing elements corresponding to an output leaf can be directly grabbed from the associated stream for output.
3. *GetQNode* performs specific edge test (PC or AD), which renders both I/O and CPU optimality for both AD and PC edges involved in a B-twig.
4. "Stack cleaning" is needed in *BTwigMerge* solely because each time after outputting path solutions, some elements on the stacks may become irrelevant for future path solutions and must be cleaned out. (In most prior algorithms such as *TwigStack*, stack cleaning is required

to get rid of those noncontributing elements that may have been tentatively added to the stacks but are actually noncontributing.)

5. *BTwigMerge* does not explicitly (at the main algorithm level) deal with any AND/OR/NOT predicates, nor with any no output *QNodes*.

#### 4.2 GetQNode: The Key Supporting Algorithm

*GetQNode* is an essential subroutine which is called by the main algorithm *BTwigMerge* to decide the next *QNode* for processing. It is *GetQNode* that guarantees that the streamhead element associated to the returned *QNode* is part of the final output since all the relevant predicates (if any) are thoroughly checked by *GetQNode* or its lower level primitive subroutines such as *edgeTest*, *nEdgeTest*, *ONodeTest*, and *hasExtension*, etc.

While feeding *BTwigMerge* with the next *QNode* to be processed, some elements on the stream under consideration may be found noncontributing to the final answer and thus should be skipped right away. The term, largest threshold value, introduced by Jiang et al. [7] refers to the start label of a sub element  $e_{\max}$  of another element, say,  $e$  such that  $e_{\max}$  maximizes the start label among all the offspring elements of  $e$ . Such a threshold value can be used to skip  $e$  and all its successors if their end label are smaller than this threshold value. It still makes sense to carry out this type of optimization for B-twig join, but we need to redefine the mechanism to fit the particular need of B-twigs.

FUNCTION *GetQNode*( $q$ )

- 1: if *isLeaf*( $q$ ) then
- 2:   return  $q$
- 3: for each  $q_i \in \text{children}(q)$  do
- 4:    $q_0 = \text{GetQNode}(q_i)$
- 5:     if  $q_0 \neq q_i$  and *isOutNode*( $q_0$ ) then
- 6:       return  $q_0$
- 7: end for
- 8:  $q_{\max} = \text{getMaxQChild}(q)$
- 9: while  $Cq \rightarrow \text{start} < Cq_{\max} \rightarrow \text{start}$  do
- 10:  $Cq \rightarrow \text{advance}()$
- 11: end while
- 12:  $q_{\min} = \text{argmin}_{q_i} \{ Cq \rightarrow \text{start} \}, q_i \in \text{children}(q)$
- 13: while  $Cq \rightarrow \text{start} < Cq_{\min} \rightarrow \text{start}$  do

```

14: if hasExtension(q) and isOutNode(q) then
15:     return q
16: else
17:     Cq → advance()
18: end while
19: if hasExtension(qmin) and isOutNode(qmin)
20: return qmin
21: else
22: Cqmin → advance()
23: if end(q) then
24: return null
25: else
26: return GetQNode(q)

```

Fig.7. Function *GetQNode*

The largest threshold value is computed by a special supporting function, called *ORBlockMax*(*n*), in [7]. We extend this function for our purpose as shown in Fig.6, which conforms to our revised notion for OR-blocks.

Understanding the structural features of OR-blocks in normalized B-twigs is the key to understanding how our *ORBlockMax* function works. This algorithm traverses the structure of an OR-block and computes the maximum threshold value to help effectively skip disqualified elements in the parent stream. Line 1 initializes the variable *q0* to a special (imaginary) query node, denoted by 0, which is always associated to a special (imaginary) element identified by the region code (0; 0; 0). When the input node is an *NQNode*, line 3 returns this special query node 0 (associated to the imaginary element (0; 0; 0)). Variable *q0* is reinitialized at line 8 to *n*, and is used at line 16 when choosing the *qmax* from all the *QNodes* *qi* under consideration such that *qmax* gives the maximal start value. At line 13, function  $\text{argmin}_{qi}(e_i.\text{start})$  selects *qmin* from all the *QNodes* *qi* under consideration such that *qmin* has the minimal start value. Notice that at this point (line 13), the imaginary element with region code (0, 0, 0) is excluded because all *NQNodes* are irrelevant to the purpose of function *ORBlockMax*—i.e., to help skip disqualified elements in the parent stream. The implementation of function *GetQNode* is shown in Fig.7. The *QNode* *q* returned by *GetQNode*(*q*) can be one of the following two cases:

**FUNCTION** *getMaxQChild*(*q*)

```

1: for each  $n_i \in \text{children}(q)$  do
2:   if  $\text{isQNode}(n_i)$ 
3:      $q_i = n_i$ 
4:   else if  $\text{isNQNode}(n_i)$  then
5:      $q_i = 0$ 
6:   else
7:      $q_i = \text{ORBlockMax}(n_i)$ 
8: end for
9: return  $\text{argmax}_{q_i} \{Cq \rightarrow \text{start}\}$ , for  $q_i$ 

```

Fig.8. Function *getMaxQChild*.

1)  $q_x = \text{null}$  (here null denotes an invalid query node), signifying to the main algorithm to immediately start another call to *GetQNode* for quickly getting the next valid *QNode* if the streams are not exhausted yet; 2)  $q_x$  is a valid output *QNode*—this is the dominating case, similarly handled as in all other holistic twig join algorithms. Comparing with *GTwigMerge* [7], the most related holistic join algorithm to *BTwigMerge*, the structure of our main algorithm is more succinct: we pushed all important tests—including AD and PC edge tests, and tests on any AND/OR/NOT predicate—all down to the core subroutine, *GetQNode*, or its lower level primitive supporting functions. The advantage is early skipping of disqualified elements in streams, leading to improved algorithm performance.

In subroutine *GetQNode(q)*, the information provided by *getMaxQChild(q)* (line 8 in Fig.7) is used to skip disqualified elements in stream  $T_q$ . Unlike its counterpart in *GTwigMerge* [7], our *getMaxQChild(q)* (see Fig.8) considers *NQNodes* in addition which do not exist in the simpler AND/OR-twigs that *GTwigMerge* was designed for.

### 4.3 Cost Analysis of *BTwigMerge*

We now analyze the I/O and CPU cost of our algorithm *BTwigMerge*. For ease of presentation, given B-twig query  $Q$ , we first introduce the following parameters:

- $[QNodes]$  is the total number of *QNodes* in  $Q$ .
- $[NQNodes]$  is the total number of *NQNodes* in  $Q$ .
- Query size  $Q_j = [QNodes] + [NQNodes]$ . Notice here we do not count other logical predicate nodes toward the query size.

- [Input] stands for the total size of all the inputstreams relevant to query Q.
- [List] stands for the average stream length.
- [Output] stands for the total count of the dataelements included in all output B-twig instancesproduced for query Q.

In terms of the set of twig patterns that can be processed, *BTwigMerge* is a “superset” of *GTwigMerge* and *GTwigMerge* is a “superset” of *TwigStack*. At the main algorithm level, the three algorithms share great similarity. The cost analysis methods are also similar. So, in the following, we only provide a compact analysis for the I/O and CPU cost of *BTwigMerge*.

The I/O cost of *BTwigMerge* consists of three parts: the I/O cost for accessing all the relevant input stream elements and the I/O cost for dealing with the intermediate path solutions plus the I/O cost for outputting the final twig solutions. Since in *BTwigMerge*, we always advance the stream cursors and never backtrack, the first part of the I/O cost is the total size of all relevant input streams. For the second part, since *BTwigMerge* is optimal with both AD and PC edges—i.e., it never produces useless intermediate path solutions, the I/O cost of this part is two times (for first output and then input) of the total final output size, i.e.,  $2 \cdot [\text{Output}]$ . And the third part (for outputting the final results), of course, is  $[\text{Output}]$ . All together, the total I/O cost for *BTwigMerge* is the sum of the above three parts. We therefore have the following equations regarding the I/O cost of *BTwigMerge*:

$$\begin{aligned} \text{I/O cost} &= ([Q\text{Nodes}] + [NQ\text{Nodes}] + [ZQ\text{Nodes}]) \cdot \\ &[\text{List}] + 3 \cdot [\text{Output}] \\ &= [Q] \cdot [\text{List}] + 3 \cdot [\text{Output}] \\ &= [\text{Input}] + 3 \cdot [\text{Output}]: \end{aligned}$$

The CPU cost analysis for *BTwigMerge* is analogous. The CPU cost also consists of three parts. The first part is the time spent on computing the path solutions, the second part is the time spent on dealing with the obtained intermediate path solutions (output, input, and merging), and the third part is on outputting the final twig solutions. The main structure of *BTwigMerge* is a loop that repeats no more than [Input] times, which is the total number of elements in all the input streams because noncontributing elements are skipped at line 10, 17, and 22 of *GetQNode* (see Fig. 13) or by the optimization rendered by the two primitive functions, *edgeTest* and *nEdgeTest* (see Figs. 8 and 9, respectively). So the first part of the CPU cost is linear to the input size. The second part depends on how many intermediate path solutions are produced and how many of



them are going to be merged to form the final output twig solutions. As *BTwigMerge* does not produce any unused intermediate path solutions (it actually does not push any noncontributing elements onto any stack), the second part of the cost is linear to and solely decided by the output size [Output]. And the third part of course is also linear to the output size. Added together, for the overall CPU cost of *BTwigMerge*, we have exactly the same result as we have for the I/O cost (cost equations omitted). It is worth to point out that query size [Q] in CPU cost is counted slightly differently from that in I/O cost: for the former, [Q] counts the duplicated query nodes caused by normalization, but for I/O cost, it does not because duplicated query nodes do not cause extra physical I/O. The above cost analysis results show that our *BTwigMerge* has both optimal I/O cost and optimal CPU cost for normalized B-twigs with both AD and PC edges. Our experimental study provides empirical evidence to further support this conclusion.

## 5 EXPERIMENTS

In this section, we present the experiment results. As our *BTwigMerge* is the only algorithm of its kind—designed for holistic B-twig pattern matching, there does not exist a real competitor to compare with. In this case, one plausible baseline to compare with is a decomposition-based approach. A decomposition-based approach first splits an input B-twig at every predicate node into a series of subtwigs, then separately computes the partial solutions to the subtwigs, and finally combines the obtained partial solutions to form the whole solutions to the original B-twig. Such a decomposition-based approach suffers severe performance disadvantage that Jiang et al. [7] had empirically proven with a subclass of B-twigs years ago. For more general B-twigs, the problem of decomposition-based approach can only become worse. So we have no intention to empirically reprove the conclusion of Jiang et al. [7] at the scale of full B-twigs; instead, we comparatively study the performance of our algorithm and other related algorithms on various common subclasses of B-twigs.

As the first holistic twig join algorithm, *TwigStack* [5] is designed for simple AND-only twigs. In terms of the categories of twigs being processed, *GTwigMerge* [7] generalizes *TwigStack* and is a superset of *TwigStack*—capable for AND/OR-twigs; *TwigStackList*: [13] also generalizes *TwigStack* but from a different aspect and thus is a superset of *TwigStack* [5] as well—capable for AND/NOT-twigs; our *BTwigMerge* significantly extends the approach embodied in *GTwigMerge* and becomes a superset of both *GTwigMerge* and *TwigStackList*:—capable for

full B-twigs, i.e., AND/OR/NOT-twigs. The theme of our experimental study thus is set on comparing *BTwigMerge*, respectively, with these predecessor algorithms with regard to a common subset of twig queries that they are all (or both) capable of dealing with.

### 5.1 Experimental Setup

Before proceeding to the details of our experiment study, we first address a few related issues about this study. Platform setup. The platform of our experiments contains an Intel Core 2 DUO 2.2 GHz running Windows XP System with 4 GB memory and a 75 GB hard disk. JavaSE is the software platform on which these algorithms are implemented and tested. The various data sets used for this study are kept as external files on the hard disk. Convenient platform, JUnit 1.4 was used for concise timing of these algorithms on test queries.

**Data preparation.** To avoid potential bias of using a single data set, we choose three popular XML data sets for this study. The first data set is an XMark data set [3] stored in a single XML file. This data set takes roughly 100 MB, containing about 100 thousands elements (or nodes). The second data set is a generated one by Stylus XML Generator [1] using a given XML Schema. Stylus XML Generator allows users to specify the expected structure and size of a XML data via separate XML Schema files. For this purpose, we carefully designed an XML schema with varied tree structures to avoid biased results.

## 6. SUMMARY

Holistic twig joins are critical operations for XML queries. The three basic logical predicates, AND, OR, and NOT, are natural expression mechanisms that people would desire to apply to general XML queries. However, all previously proposed holistic twig join algorithms failed to provide an integral solution for efficient and uniform processing of B-twig queries (with arbitrary combination of these logical predicates) in a single algorithmic framework. Consequently, given a general B-twig query, all prior holistic algorithms become inapplicable and useless. In this paper, we presented a novel approach for holistic computing of B-twig patterns and described an original algorithm, called *BTwigMerge*, which is the first of its kind—holistic computing of a more general class of twig patterns represented by B-twigs. The second distinctive feature of *BTwigMerge* is that it gracefully extends the I/O and CPU optimality to twigs with PC edges as well.

In order to reduce the intrinsic complexity in arbitrary B-twigs, we proposed B-twig normalization that successfully sorts out the arbitrary combination of the logical predicates in B-twigs. We designed a valid procedure to automatically transform input B-twigs into normalized forms. The normalized B-twigs are then sent to BTwigMerge that embodies our holistic twig join strategy and contains numerous novel supporting mechanisms.

We have done analytical and experimental study with regard to the validity and performance of our approach and its accompanying algorithms, and concluded that our BTwigMerge is so far the most powerful and most efficient holistic twig join algorithm—the sole one designed for B-twigs, with optimal I/O and optimal CPU on twigs with arbitrary AD and/or PC edges. As future work, the following is on our agenda: Implementation of XOR and efficiency increase by Indexing method.

## REFERENCES

- [1] Stylus Studio XML Generator, [http://www.stylusstudio.com/xml\\_generator.html](http://www.stylusstudio.com/xml_generator.html), 2012.
- [2] Univ. of Washington XML Repository, <http://www.cs.washington.edu/research/xml/datasets/>, 2012.
- [3] XMark? An XML Benchmark Project, <http://www.xmlbenchmark.org/>, 2012.
- [4] S. Al-Khalifa et al., “Structural Joins: A Primitive for Efficient XML Query Pattern Matching,” Proc. 18th Int’l Conf. Data Eng. Conf. (ICDE’02), pp. 141-152, 2002.
- [5] N. Bruno, N. Koudas, and D. Srivastava, “Holistic Twig Joins: Optimal XML Pattern Matching,” Proc. ACM SIGMOD Int’l Conf. Management of Data (SIGMOD’02), pp. 310-321, June 2002.
- [6] T. Chen, J. Lu, and T.W. Ling, “On Boosting Holism in XML Twig Pattern Matching Using Structural Indexing Techniques,” Proc. ACM SIGMOD Int’l Conf. Management of Data (SIGMOD’05), pp. 455-466, June 2005.
- [7] H. Jiang, H. Lu, and W. Wang, “Efficient Processing of Twig Queries with OR-Predicates,” Proc. ACM SIGMOD Int’l Conf. Management of Data (SIGMOD’04), pp. 59-70, 2004.
- [8] H. Jiang, W. Wang, H. Lu, and J.X. Yu, “Holistic Twig Joins on Indexed XML Documents,” Proc. 29th Int’l Conf. Very Large Data Bases (VLDB’03), pp. 273-284, Sept. 2003.

- [9] J. Lu, T. Chen, and T.W. Ling, "Efficient Processing of XML TwigPatterns with Parent Child Edges: A Look-ahead Approach," Proc.13th ACM Int'l Conf. Information and Knowledge Management(CIKM' 04), pp. 533-542, Nov. 2004.
- [10] J. Lu, T.W. Ling, C.-Y. Chan, and T. Chen, "From Region Encoding to Extended Dewey: On Efficient Processing of XML Twig Pattern Matching," Proc. 31st Int'l Conf. Very Large Data Bases (VLDB' 05), pp. 193-204, Aug. 2005.
- [11] J. Lu et al., "Efficient Processing of Ordered XML Twig Pattern," Proc. 16th Int'l Conf. Database and Expert Systems Applications(DEXA' 05), pp. 300-309, 2005.
- [12] X. Xu, Y. Feng, and F. Wang, "Efficient Processing of XML Twig Queries with All Predicates," Proc. IEEE/ACIS Int'l Conf. Computer and Information Science (ICIS ' 09), pp. 457-462, June 2009.
- [13] T. Yu, T.W. Ling, and J. Lu, "twigstacklist: A Holistic Twig Join Algorithm for Twig Query with Not-Predicates on XML Data," Proc. 11th Int'l Conf. Database Systems for Advanced Applications (DASFAA' 06), pp. 249-263, 2006.
- [14] XML Repository <http://ghr.nlm.nih.gov/search>