Pragmatic XML Access Control Using Off-the-Shelf RDBMS

Bo Luo, Dongwon Lee, and Peng Liu*

The Pennsylvania State University
{b1uo,dongwon,px120}@psu.edu

Abstract. As the XML model gets more popular, new needs arise to specify *access control* within XML model. Various XML access control models and enforcement methods have been proposed recently. However, by and large, these approaches either assume the support of security features from XML databases or use proprietary tools outside of databases. Since there are currently few commercial XML databases with such capabilities, the proposed approaches are not yet practical. Therefore, we explore the problem of “Is it possible to fully support XML access control in RDBMS?” We formalize XML and relational access control models using *deep set* operators. Then we show that the problem of XML AC atop RDBMS is amount to the problem of converting XML deep set operators into *equivalent* relational deep set operators. We show the conversion algebra and identify the properties to ensure the correct conversion. Finally, we present three practical implementations of XML access controls using off-the-shelf RDBMS and their performance results.

1 Introduction

The XML model [1] has emerged as the *de facto* standard for storing and exchanging information in the Internet Age. As more information is exchanged over the Web, the issues of security become increasingly important. Such issues span from data level security to network level security to high-level access controls. In this paper, our focus is on how to support *access control* for XML data. Many access control methods extending the XML model to incorporate security aspects have been proposed recently (e.g., [2,3,4,5]). To the lesser or greater extent, however, XML access control enforcement mechanisms proposed in the research community neglect the fact that the most XML data still resides in RDBMS. In the scenario of RDBMS-backed XML database systems (hereafter *XRDB*), XML documents are stored in RDBMS and query-answering is conducted through a conversion layer. In the scenario of XML publishing, relational data is compiled into XML format for distribution and exchange. For both scenarios, we enjoy the benefit of XML model while taking advantage of the maturity of the off-the-shelf RDBMS. In both scenarios, it is desirable to natively specify access controls on the XML side, but they need to be enforced on the

* Peng Liu was supported in part by NSF CCR-0233324.

RDBMS. We believe that current XML access control enforcement mechanism research is in a sense re-inventing wheels without utilizing existing relational access control models or leveraging on security features that are readily available in relational products. In short, our goal in this paper is to answer the question: *When is it (not) possible to support XML access control using RDBMS? Why? How?*

Challenges. First, the major challenges of supporting XML access control in XRDB systems stem from the inherent discrepancy of XML and relational data models. Relational data model features a structure of two-dimensional table, while XML features a hierarchical data model. When XML data are shredded into relational data model, not all transformation algorithms can fully preserve structural properties of XML model [6]. The inherent incompatibility of two data models leads to the fundamental discrepancy between two access control models. Second, relational access control policies define authorized actions of “cells,” where each cell is an impartible element and whose accessibility is explicitly expressed. However, XML nodes are hierarchically nested, and XML data model inherently takes “answer by subtree model” (e.g., querying for `//foo` yields the whole subtree rooting at node `<foo/>`). Therefore, for any XML node, an action could be: authorized/unauthorized to the whole subtree, or partially authorized. The later case does not occur in relational access control model. Finally, in XML model, we can control the access right of each individual node. In traditional relational model, the smallest granularity that we may control is a column via GRANT/REVOKE. Therefore, we need to employ more recent developments of RDBMS access controls (e.g., Oracle VPD) to enable cell level access control.

Key contributions. (1) To our best knowledge, this work is the first one to algebraically formalize XML access control in both native XML (XDB) and XRDB environment. (2) This work takes the first steps to define the *equivalent objects* and *equivalent operations* between native XML and XRDB systems. With this concept, we can migrate all the exciting features of native XML systems into XRDB by converting the atomic operations into equivalent relational counterparts. In this paper, we take the feature of fine-grained XML access control for a pilot study, and the results are encouraging. (3) This work shows for the first time that the “security” of XRDB can be achieved by finding the “equivalent” relational operators for three specific deep-set operators. This finding provides a viable way to build secure XRDB systems. (4) Finally, this work proposes several practical approaches to implement the viable way “discovered” by our theory.

2 Related Work

2.1 XML and Relational Access Control

Current access control research can be categorized into two groups: access control modeling and access control enforcement mechanisms.

On the model side, several XML access control models have been proposed. Starting with [7] for HTML documents; [8,3] describes XML access control with

Fig. 1. Overview architecture and an example XML schema

an authorization sheet to each document or DTD. [4] proposed an XML access control model to deal with authorization priorities and conflict resolution. [9] introduced provisional authorization and XACL. [10] formalizes the way of specifying *objects* in XML access control using XPath. Most of the proposals adopt either *role-based access control* (e.g. [12]) or *credential-based access control* (e.g. [13]). The major difference between them is the way they identify users. Credential-based access control is more flexible and powerful in this aspect. However, in the research of access control enforcement mechanisms, people tend to choose a relatively simple access control model to avoid distraction.

XML access control enforcement mechanisms in native XML environment have been intensively studied in recent years. They are categorized into four classes: (1) engine level mechanisms implement security check inside XML database engine; each XML node is tagged with a label [15,16,17] or an authorization list [18,19], and filtered during query processing. (2) view based approaches build security views that only contain access-granted data [20,21,22].¹ (3) pre-processing approaches check user queries and enforce access control rules before queries are evaluated, such as the static analysis approach [23,14], QFilter approach [24], access condition table approach [25], policy matching tree [26], secure query rewrite (SQR) approach [27], etc. (4) [28] considers access control of streaming XML data and apply security check at client side, using a filtering mechanism. More recently, [29,30] takes encryption issues into consideration, and [31,32] focus on protecting the privacy and security associated with XML tree structure.

Relational access control models can be classified into two categories: *multi-level security models* [33,34,35] and *discretionary security models (DAC)*. Most real world database systems implement a table/column level DAC similar to the one implemented in System R [36]. View-based approaches is the traditional method to enable row-level access control, while Oracle's VPD is the most recent development. Finally, some advanced access control models (e.g., [37,38]) are proposed in a more theoretical manner.

2.2 XML and Relational Conversion

As illustrated in Figure 1(a), in XRDB: XML data D_X are converted into D_R and stored in RDBMS; user issues XML query Q_X (XPath or XQuery) using

¹ When a view-based approach implements virtual views without materializing them, it is inherently a pre-processing approach.

published XML schema; Q_X is converted into Q_R (SQL) and evaluated against D_R ; relational answer A_R is finally converted to XML (A_X) to return to user.

Toward conversion between XML and relational models, an array of research has addressed the particular issues lately. On the industry side, database vendors are busily extending their databases to adopt XML types. Shredding and non-shredding are two major pathes that followed by commercial products. Oracle provides both un-shredded (CLOB) and shredded storage options [40]. Microsoft supports XML shredding and publishing through mid-tier approach in SQL Server 2000, and adds CLOB storage in SQL Server 2005 [41]. IBM proposes the first native XML storage in DB2 9, but shredded XML storage (through schema decomposition) is still kept as an important feature [42,43]. On the research side, various proposals have been made recently, mainly either schema-based (e.g., [44,45,46]) or schema-oblivious (e.g., [47,48]) approaches.

In terms of access control, some commercial products apply existing column level access control of RDBMS on XML data stored in CLOB columns. None of these approaches supports or discusses fine-grained access control. Finally, to our best knowledge, the only work that is directly relevant to our proposal is [5]. [5] proposes an idea of using RDBMS to handle XML access controls, in a rather limited setting. In our vision paper [49], we addressed some issues and challenges of enforcing XML access control atop RDBMS. We provide the algebraic analysis and explore practical solutions in this paper.

Our framework is not tied to a particular conversion method. Throughout this paper, we use shared-inlining [45] and XRel [48] as the examples of schema-based and schema-oblivious conversion methods, respectively. Briefly, XRel decomposes XML documents into document, element, attribute, text, and path tables. In this approach, each node is stored as one record in the `element` table, and each distinct path is stored as one record in the `pth` table. As a simple example, we decompose an XMark ([50], Figure 1(b)) document using XRel and show part of `element` table in Figure 2 (b). As we can see, element 252 is a node of path 164 (“/site/people”, as stored in the `path` table); which starts from offset 33996 (byte) and ends at 36229 in the original XML document.

3 Preliminaries

3.1 XML Access Control Policy

Access control models define the semantics and syntax of access control policies. Although they could be very complicated, the essential is to describe *subjects*, *objects*, *actions* and all the variations around it. Fortunately, there is no discrepancy in identifying *subjects* and defining *actions* in XML and relational environment, e.g., they both could adopt RBAC or CBAC to identify users. As we described in Section 1, shredding XML access control models into relational ones is a challenging task, because of the fundamental discrepancies of XML and Relational data models. Therefore, challenges reside in *object*-related components of access control models, while issues that only relate to *subjects* and *actions* are trivial. Thus, our subsequent discussion focuses more on *object* part.

In this paper, we adopt the model proposed in [3] as the basis; other models like [23,28,24,51] can be used as well with a reasonable change.

Definition 1 (XML Access Control Rule) . An XML access control policy is specified by a set of **access control rules**: $R_X = \{\text{subject, object, action, sign}\}$, where *subject* is to whom an authorization is granted, *object* is a set of XML nodes (in XPath) to which the policy is applied, *action* is one of read, write, or update, and $\text{sign} \in \{+, -\}$ refers to access granted or denied, respectively. $\square$

In this model, access is prohibited by default. Negative rule takes precedence when it conflicts with positive rules. All access controls propagate to the entire subtree rooting at *object*, complying with the answer-by-tree XML semantics. If a rule applies to context node only, we add “/text()” to its *object* field.

3.2 XML to Relational Conversion

Remark 1. A relational to XML conversion method contains: (1) $C_D()$ to convert XML to relational data, (2) $C_Q()$ to convert XML query (XQuery or XPath) to SQL, and (3) C^{-1} to convert relational answer back to XML. $\square$

That is, $Q_R = C_Q(Q_X)$, $D_R = C_D(D_X)$, and $A_X = C^{-1}(A_R)$. From this, the process of “evaluating XML query on XRDB” can be modeled as:

$$A_X = C^{-1}(A_R) = C^{-1}(Q_R(D_R)) = C^{-1}(C_Q(Q_X)(C_D(D_X))) \quad (1)$$

Remark 2 . An X2R conversion algorithm is lossless iff: (1) (lossless node conversion) $\forall$ XML node x_i , $C_D^{-1}(C_D(x_i)) = x_i$; (2) (lossless node set decomposition) $\forall$ XML node set $\{x_1, \dots, x_n\}$, $C_D^{-1}(C_D(\{x_1, \dots, x_n\})) = C_D^{-1}(\{C_D(x_1), \dots, C_D(x_n)\}) = \{C_D^{-1}(C_D(x_1)), \dots, C_D^{-1}(C_D(x_n))\}$; and (3) (exclusive conversion) $C_D(x_1) = C_D(x_2)$ only when $x_1 = x_2$, and $C_D^{-1}(r_1) = C_D^{-1}(r_2)$ only when $r_1 = r_2$. $\square$

Remark 3 . An X2R conversion algorithm is correct iff: $\forall$ query Q and $\forall$ document X , $Q(X) = C^{-1}(Q_R(D_R)) = C^{-1}(C_Q(Q_X)(C_D(X)))$. $\square$

Definition 2 (Soundness) . An X2R conversion algorithm A is **sound** iff it is lossless and correct. $\square$

In the remainder of the paper, we assume that the conversion algorithm being used is *sound*. We ignore the order of XML nodes when we compare the correctness, since this feature is not supported in most X2R conversion algorithms.

In the research community, most X2R conversion algorithms support a subset of XQuery/XPath (e.g., /, //, * and predicates). Our approach does not alter the query or data conversion algorithm. Therefore, for a particular X2R conversion method X , we support everything that X supports. For ease of understanding, we do not use predicates in the examples, however, we test queries with predicates in our experiments.

3.3 Deep Set Operators

In [52], we propose deep set operators for XML, as extensions of conventional set operators defined in XPath [53] and XQuery [54]. Here, we briefly revisit them, and later demonstrate how they are used to formalize XML access control.

Definition 3 (deep set operators). The **deep-union** operator ($\overset{D}{\cup}$) takes node sequences $\langle P \rangle$ and $\langle Q \rangle$ as operands, and returns a sequence of nodes (1) who exist as a node or as a descendant in **either** operand sequences, and (2) whose parent does not satisfy (1). Formally, $\langle P \rangle \overset{D}{\cup} \langle Q \rangle = \{n | (n \in \langle P_d \rangle \vee n \in \langle Q_d \rangle) \wedge (n :: \text{parent}() \notin \langle P_d \rangle \wedge n :: \text{parent}() \notin \langle Q_d \rangle)\}$ where $P_d = P/\text{descendant} - \text{or} - \text{self}()$. The **deep-intersect** operator ($\overset{D}{\cap}$) takes node sequences $\langle P \rangle$ and $\langle Q \rangle$ as operands, returns a sequence of nodes (1) who exist as a node or as a descendant in **both** operand sequences, and (2) whose parent does not satisfy (1). Formally, $\langle P \rangle \overset{D}{\cap} \langle Q \rangle = \{n | (n \in \langle P_d \rangle \wedge n \in \langle Q_d \rangle) \wedge (n :: \text{parent}() \notin \langle P_d \rangle \vee n :: \text{parent}() \notin \langle Q_d \rangle)\}$. Finally, the **deep-except** operator ($\overset{D}{-}$) takes node sequences $\langle P \rangle$ and $\langle Q \rangle$ as operands, for each node $\langle p_i \rangle$ in $\langle P \rangle$, it remove $\langle p_i \rangle \overset{D}{\cap} \langle Q \rangle$ from the subtree of $\langle p_i \rangle$ and return the remaining. $\square$

3.4 XML Access Control in XDB and XRDB

XML access control is to ensure that only *safe answer* (SA) is returned. As in [55,52], safe answer of Q includes all the XML nodes n such that: (1) $n \in \langle Q \rangle$, (2) the access of n is granted by positive rules, and (3) the access of n is *not* denied by negative rules. Therefore, the precise semantics of “safe answer,” SA_X is:

$$SA_X = \langle Q_X \rangle \overset{D}{\cap}_X [(\langle R_{X_1}^+ \rangle \overset{D}{\cup}_X \dots \overset{D}{\cup}_X \langle R_{X_n}^+ \rangle) \overset{D}{-}_X (\langle R_{X_1}^- \rangle \overset{D}{\cup}_X \dots \overset{D}{\cup}_X \langle R_{X_m}^- \rangle)] \quad (2)$$

Equation (1) models XML query evaluation in XRDB. Similarly, (2) models how only safe XML answers, SA_X , are returned. Combine them, we have:

Definition 4 (Secure XRDB). An XRDB is secure iff $\forall$ ACR set ACR_X and $\forall$ query Q_X , it always returns the safe answer: $A_X \equiv SA_X$. Therefore,

$$C^{-1}(\{C_Q(Q_X)\langle C_D(D_X) \rangle\}') \equiv \langle Q_X \rangle \overset{D}{\cap}_X [(\langle R_{X_1}^+ \rangle \overset{D}{\cup}_X \dots) \overset{D}{-}_X (\langle R_{X_1}^- \rangle \overset{D}{\cup}_X \dots)] \quad (3) \quad \square$$

Note that $\{C_Q(Q_X)\langle C_D(D_X) \rangle\}'$ indicates that access control mechanism intervenes in relational query processing. Our goal in this paper is to enforce XML access controls on RDBMS so that Equation 3 holds in XRDB setting. In this way, we need to convert *access control rules* R_X and *deep set operators* into their *equivalent* relational counterpart.

1. {user, /site/people/person, read, +}					SELECT e0.DOCID, e0.ELEMENTID, e0.PATHID, e0.ST, e0.ED				
2. {user, /site/people/person/credicard, read, -}					FROM document d, element e0, pth p0				
(a)					WHERE p0.pathexp LIKE '%/people'				
					AND e0.pathid = p0.pathid AND d.docid = e0.docid				
DOCID	ELEMENTID	PATHID	ST	ED	(c)				
0	252	164	33996	36229	<people>	SELECT e0.DOCID, e0.ELEMENTID, e0.PATHID, e0.ST, e0.ED			
0	293	165	35592	35826	<person>	FROM document d, element e0, pth p0			
0	299	165	35832	36217	<person>	WHERE p0.pathexp LIKE '%/people//person'			
0	303	188	35989	36032	<creditcard>	AND e0.pathid = p0.pathid AND d.docid = e0.docid			
(b)					(d)				

Fig. 2. Naive enforcement of “equivalent” relational ACR leads to wrong answer

4 XML Access Control in XRDB: The Theory

All entities of the 4-tuple XML access control model, except *object*, can be directly adopted to relational access control model. We apply an X2R algorithm $C(R_X.object)$ to get $R_R.object$. Therefore, we get “equivalent” relational ACR:

$$\mathbf{R}_R = \{\mathbf{R}_X.subject, C(\mathbf{R}_X.object), \mathbf{R}_X.action, \mathbf{R}_X.sign\}$$

However, naive enforcement of the converted relational access control rules may lead to security leakage, as demonstrated in the following example:

Example 1. Consider two rules of Fig. 2(a) with XRDB(XRel) – an XRDB employing XRel [48] as the conversion algorithm. The “element” table is partly shown in Fig. 2(b). Rule 1 indicates that a user is allowed to access **<person>** nodes, i.e., nodes 293 and 299 (record 2 and 3 in Fig. 2 (b)), and rule 2 indicates that a user cannot access **<credicard>** nodes, i.e., node 303. Naive enforcement will grant access to the record 2, 3; and revoke the access to record 4.

Query “//people” is desired to yield an answer containing two **<person>** nodes, since they are the accessible descendants of the requested node. However, the converted SQL query (Fig. 2(c)) yields no answer since access to record 1 is prohibited by default. Moreover, for a query “//person”, the converted SQL (Fig. 2(d)) returns both **<person>** nodes to the user (with the unauthorized **<creditcard>** node). This is so because both records of element 293 and 299 are accessible, while revoking access to element 303 does not affect its ancestor. □

4.1 Object and Operation Equivalency

To solve the problem illustrated in Example 1, we propose our framework of supporting access control in XRDB systems. First, we define object and operation equivalency between XML and relational.

Definition 5 (Object Equivalency). When both $R = C(X)$ and $X = C^{-1}(R)$ hold for XML node set X and relation R , we consider X and R equivalent w.r.t. C/C^{-1} , and denote as $X \equiv R$. □

Note that, when we talk about equivalency of X and R , we have to predefine the context, i.e., select the X2R conversion algorithm C/C^{-1} . For a XML node set X , $C(X)$ may be different under different X2R conversion algorithms.

Definition 6 (Operation Equivalency). Suppose $X_1 \equiv R_1$ and $X_2 \equiv R_2$ w.r.t. C/C^{-1} . Then, an XML operation OP_X is equivalent to a relational operation OP_R (denoted as $OP_X \equiv OP_R$) w.r.t. C and C^{-1} if:

$$C(X_1 \text{ } OP_X \text{ } X_2) = C(X_1) \text{ } OP_R \text{ } C(X_2) = R_1 \text{ } OP_R \text{ } R_2 \quad \square$$

Note that XML operator takes node sets as operands while its equivalent relational counterpart may not take two generic relations as operands. Each operand is the equivalent objects of corresponding XML node set, which may be tables, columns, records, etc. Relational operations require operands to be domain compatible (e.g., intersect, union etc.). We loosen this requirement for OP_R .

With the concept of operation equivalency, we can migrate all the exciting features of XML into XRDB by converting the atomic operations into equivalent relational operation. Our secure XRDB problem is articulated as follows:

Lemma 1. In $XRDB(C)$, if we can find relational operators, $\overset{D}{\cup}_R$, $\overset{D}{\cap}_R$, and $\overset{D}{-}_R$, which are equivalent to XML deep set operators, $\overset{D}{\cup}_X$, $\overset{D}{\cap}_X$, and $\overset{D}{-}_X$, w.r.t. the X2R conversion algorithm C , we are able to enforce XML access control in $XRDB(C)$ such that Equation (3) always holds. ■

Please refer to [56] for detailed proof. Now we need to find equivalent operations such that $\overset{D}{\cup}_R \equiv \overset{D}{\cup}_X$, $\overset{D}{\cap}_R \equiv \overset{D}{\cap}_X$ and $\overset{D}{-}_R \equiv \overset{D}{-}_X$. Again, equivalency is based on specific X2R conversion method, therefore, the existence and representation of relational deep set operators also depends on the particular X2R conversion. Hereafter, we analyze the role of each deep set operator in (2) and the existence of its equivalent relational counterpart under different X2R conversion algorithms.

4.2 On Equivalent Conversion of Deep Set Operators

Deep-union operator is used to integrate all the accessible nodes defined by individual positive rules (also, all the inaccessible nodes defined by negative rules). With the property $P \overset{D}{\cup} Q \subseteq P \cup Q$ [52], Remark 1 is rewritten into:

$$\langle P \rangle \overset{D}{\cup}_X \langle Q \rangle = \{n | (n \in \langle P \rangle \vee n \in \langle Q \rangle) \wedge (n \notin \langle P//* \rangle \wedge n \notin \langle Q//* \rangle)\} \quad (4)$$

Let $r = C(n)$. When C/C^{-1} is sound according to Definition 2, we have:

$$C(\langle P \rangle \overset{D}{\cup}_X \langle Q \rangle) = \{r | [r \in C(\langle P \rangle) \vee r \in C(\langle Q \rangle)] \wedge [r \notin C(\langle P//* \rangle) \wedge r \notin C(\langle Q//* \rangle)]\}$$

Here, since we are to find $\overset{D}{\cup}_R$ such that $C(\langle P \rangle) \overset{D}{\cup}_R C(\langle Q \rangle) = C(\langle P \rangle \overset{D}{\cup}_X \langle Q \rangle)$:

$$C(\langle P \rangle) \overset{D}{\cup}_R C(\langle Q \rangle) = \{r | [r \in C(\langle P \rangle) \vee r \in C(\langle Q \rangle)] \wedge [r \notin C(\langle P//* \rangle) \wedge r \notin C(\langle Q//* \rangle)]\}$$

The condition of $[r \in C(\langle P \rangle) \vee r \in C(\langle Q \rangle)]$ is essentially the regular union. It is composed by set containment and Boolean operations. In XRDB, set containment check is supported when the soundness requirement in Definition 2 is fulfilled, and Boolean operation is generally supported in RDBMS. $[r \notin C(\langle P//* \rangle)]$

$\wedge r \notin C(\langle Q//*\rangle)]$ tends to support deep semantics. It requires XRDB to be able to identify if $r \in C(\langle P//*\rangle)$ for any given relational object r and set $C(\langle P\rangle)$.

Lemma 2. *To implement deep-union operator in XRDB(C), the X2R conversion algorithm C should: (1) fulfil the soundness requirement stated in Definition 2; and (2) for given node n and node set $\langle P\rangle$, it should be able to check the containment condition of: $C(n) \in C(\langle P//*\rangle)$, e.g., it should recognize if $C(n)$ is a descendant of any node $C(p_i)$; ■*

At present, all X2R conversion algorithms (we are aware of) fulfill Lemma 2.

Deep-intersect operator is used to calculate the exact overlapping of queried data and accessible data (i.e. $\langle Q\rangle$ and $\langle ACR\rangle$). It is defined as:

$$C(\langle P\rangle \overset{D}{\cap}_X \langle Q\rangle) = \{r \mid [r \in C(\langle P\rangle) \wedge r \in C(\langle Q\rangle)] \wedge [r \notin C(\langle P//*\rangle) \vee r \notin C(\langle Q//*\rangle)]\} \quad (5)$$

Compare with 4, they only differ in logical operators. Therefore, Lemma 2 could be directly extended to deep-intersect.

Example 2. In Example 1, a query “//people” yields $\langle \text{people} \rangle$ nodes, i.e. element 252, (record 1 in Fig. 2 (b)). Meanwhile, *object* field of access control rule 1, “/site/people/person”, yields $\langle \text{person} \rangle$ nodes, i.e. element 293 and 299 (record 2 and 3 in Fig. 2 (b)). In XRel, each XML node is marked with “start” and “end” offset. Node containment is checked through comparison of the offsets: for nodes p_1 and p_2 , if $(p_1.start < p_2.start)$ and $(p_1.end > p_2.end)$, p_2 is an descendant of p_1 . Here, we can tell that node 293 and 299 are descendants of node 292. Therefore, “//people $\overset{D}{\cap}_X$ //person” yields node 293 and 299. Comparing with Example 1, “//people $\cap$ //person” yields *Null*. □

The operands of XML deep-union/intersect operators may contain different nodes. In RDBMS, where domain compatibility is strictly enforced, their relational equivalent counterpart might be domain incompatible (e.g. a row “intersect” a cell). This happens in schema-based X2R conversion (e.g. [44,45]), where different XML nodes could be converted to tables, rows, etc. To tackle this problem, we employ new RDBMS techniques, e.g. Oracle VPD, to enable fine-control of relational tables to create relational views with any group of cells.

Deep-except is used to remove inaccessible nodes from the answer. Recall that, in our XML access control model, all nodes are inaccessible by default. When a user is prohibited to access a node, there is no need to write a negative rule to revoke accessibility unless the node is covered by positive rules (ACR^+). Thus, negative rules are only used to specify exceptions to global permissions, i.e. “revoke” access granted by ACR^+ . Deep except operator is used to enforce negative rules. Regarding whether deep except could be implemented in XRDB with X2R conversion algorithm, it depends upon the characteristics of specific negative rule. In particular, we distinguish two types of negative rules:

Definition 7 (Node elimination vs. Descendant elimination rules). A negative rule in ACR restricts user from access a set of nodes $\{r_1^-, \dots, r_n^-\}$. If **none** of the nodes is a descendant of the context node of a positive rule, i.e.:

$$r_i^- \notin \langle R^+ // * \rangle, \quad \forall r_i^- \in \{r_1^-, \dots, r_n^-\}; \forall \langle R^+ \rangle \in \langle ACR^+ \rangle$$

then it is called a **node elimination (NE)** negative rule. Else, if one of the nodes is a descendant of the context node of a positive rule, i.e.:

$$r_i^- \in \langle R^+ // * \rangle, \quad \exists r_i^- \in \{r_1^-, \dots, r_n^-\}; \exists \langle R^+ \rangle \in \langle ACR^+ \rangle$$

it is called a **descendant elimination (DE)** negative rule. □

Intuitively, “Node Elimination” negative rule removes context node from $\langle ACR^+ \rangle$. For XML nodes covered by node elimination negative rules $\langle ACR_1^- \rangle$, deep-except operator directly removes them from $\langle ACR^+ \rangle$:

$$\langle ACR^+ \rangle \stackrel{D}{-}_X \langle ACR_1^- \rangle = \{n | n \in \langle ACR^+ \rangle \wedge n \notin \langle ACR_1^- \rangle\}$$

Essentially, this is the regular except semantics. In this way, in XRDB, we have,

$$C(\langle ACR^+ \rangle) \stackrel{D}{-}_R C(\langle ACR_1^- \rangle) = \{r | r \in C(\langle ACR^+ \rangle) \wedge r \notin C(\langle ACR_1^- \rangle)\}$$

To support deep except operator for node elimination negative rules only, the conditions described in Lemma 2 still apply.

On the other hand, “Descendant Elimination” negative rule removes descendants from context node of $\langle ACR^+ \rangle$. It takes more burden to process descendant elimination negative rules, where real “deep” semantics is required. That is,

$$\langle ACR^+ \rangle \stackrel{D}{-}_X \langle ACR_2^- \rangle = \{deepRemove(n, n \bigcap_X \langle ACR_2^- \rangle) | n \in \langle ACR^+ \rangle\}$$

where $deepRemove(p, \langle Q \rangle)$ takes a node and a set of its descendants as operands, removes the descendants from the subtree of the node and return the remaining. This function may not be directly converted to relational.

Lemma 3. When deep-except operator takes node specified by descendant elimination negative rules as the second operand, it is implemented through $deepRemove()$ operation. To implement deep-except operator that supports descendant elimination negative rules in $XRDB(C)$, the $X2R$ conversion algorithm X should: (1) fully satisfy Lemma 2; and (2) for any node n_1 and its descendant n_2 , $C(n_2)$ should be part of $C(n_1)$; and in the reverse conversion of $n_1 = C^{-1}(C(n_1))$, node n_2 in the subtree is entirely converted from $C(n_2)$. ■

Example 3. For instance, in Example 1, Rule 2 is a descendant elimination negative rule since it revoke access towards descendants of node (**<person>**).

In **XRDB(XRel)** [48], descendants are converted to independent records that are stand alone from ancestors. As shown in Fig. 2(b), to reconstruct a **<person>** node, $C_{XRel}^{-1}()$ only takes record 2 to reconstruct a full node. The descendant

`<creditcard>` is included in the answer, but the record 4 is not touched by $C_{XRel}^{-1}()$. Therefore, XRel violates condition (2) of Lemma 3, so that we cannot directly implement deep-except operator to support descendant elimination rules. When user requests for `//person`, we are not able to revoke access towards `<creditcard>` child from RDBMS side.

In **Shared-Inlining** [45] approach, `<person>` nodes are translated into a table, and `<creditcard>` nodes take a column. The relational schema is [57]:

```
Person(Id, ParentId, Person, Person_address, Person_address_city, ...
..., Person_address_zipcode, Person_creditcard, .....)
```

Here, the ancestor-descendant relationship is kept such that each row represents a “person” node, and each cell represents a child node. Therefore, to obtain `//personXD//creditcard`, we just mask “person_creditcard” column in the table; and the reconstructed XML tree of “person” node will not have corresponding child, i.e., “creditcard” node is removed from the XML answer. $\square$

5 XML Access Control Enforcement in XRDB

Figures 3 shows a general framework for XML access control in XDB and XRDB. There are three categories of XML access control enforcement mechanisms: (1) view-based approach (① ④ in Fig. 3); (2) pre-processing approach (② ⑤ in Fig. 3); and (3) post-processing approach (③ ⑥ in Fig. 3). In this section, we articulate the algebra of these approaches using deep set operators.

5.1 View-Based Approach

When access control is first enforced on XML documents to create *views*, it is the traditional view-based approach. In this model, XML view V_X (or *safe document SD*) is constructed, and query is evaluated against the view

$$SA = Q \langle V_X \rangle = Q[(\langle R_{X1}^+ \rangle \overset{D}{\cup}_X \dots \overset{D}{\cup}_X \langle R_n^+ \rangle) \overset{D}{-}_X (\langle R_1^- \rangle \overset{D}{\cup}_X \dots \overset{D}{\cup}_X \langle R_m^- \rangle)]$$

Fig. 3. Access control enforcement approaches in XML DB and XRDB

Fig. 4. Enforcing XML access control via external pre-processing

To convert this approach into XRDB, we can either convert XML view into relational view, as shown in ① of Figure 3; or construct relational view using converted relational ACR, ④ of Figure 3. They are formalized as:

$$SA = C^{-1}(Q_R \langle V_R \rangle) = C^{-1}(Q_R \langle C(V_X) \rangle) \quad (6)$$

$$\text{and} \quad SA = C^{-1}(Q \langle V_R \rangle) = C^{-1}(Q \langle C(ACR_X) \langle D_R \rangle \rangle) \quad (7)$$

5.2 Pre-processing Approach

In preprocessing model, *safe query* SQ is constructed as:

$$SQ_X = Q_X \overset{D}{\cap}_X [(R_{X1}^+ \overset{D}{\cup}_X \dots \overset{D}{\cup}_X R_{Xn}^+ \overset{D}{\cap}_X (R_{X1}^- \overset{D}{\cup}_X \dots \overset{D}{\cup}_X R_{Xm}^-)]$$

Safe answer is yielded by evaluating safe query against the original document: $SA_X = SQ_X \langle D_X \rangle$. To extend this approach to XRDB, we have: (1) XML Query Rewriting: as shown in ② in Fig. 3, we convert the safe XML query into SQL, and answer it with regular XRDB; and (2) Relational Query Rewriting: as shown in ⑤ in Fig. 3, we convert original Q_X into SQL Q_R , and then we rewrite it into safe query SQ_R . They are formalized as

$$SA_X = C^{-1}(SQ_R \langle D_R \rangle) = C^{-1}(C(SQ_X) \langle D_R \rangle) \quad (8)$$

$$\text{and} \quad SQ_R = Q_R \overset{D}{\cap}_R [(R_{R1}^+ \overset{D}{\cup}_R \dots \overset{D}{\cup}_R R_{Rn}^+ \overset{D}{\cap}_R (R_{R1}^- \overset{D}{\cup}_R \dots \overset{D}{\cup}_R R_{Rm}^-)] \quad (9)$$

Example 4. Let us revisit the previous examples: we manage XMark document in XRDB(XRel). Suppose we have access control rule (user, //people, read, +), and user submits query //name. Figure 4(a) shows the relational query for $C(//people) \overset{D}{\cap}_R C(//name)$, which is implemented according to the definition in Equation 5 (we marked up all the sub-queries). Moreover, this query could be further optimized, as shown in Figure 4(b). $\square$

Another method is to use Oracle VPD. Oracle version 8.1.5 introduces a new security feature supporting non-view-based fine-grained access control, namely *Row Level Security* or *Virtual Private Database*. It allows users to control accessibility towards row/cell level. With VPD, we are able to tailor relational data into any shape we want. To utilize VPD for access control in XRDB, we first construct relational predicates from the converted relational access control rules ACR_R , then define a VPD policy to enforce the predicates on converted SQL queries. Moreover, cell level access control capability of VPD is of special importance to XRDB systems that use schema-based X2R conversion algorithm, such as Inlining. In those XRDB systems, XML nodes are converted to different types of relational objects: tables, rows and cells. In this way, $\langle ACR \rangle$ may not be conventional relations, e.g. it could be arbitrary combinations of columns, rows and/or individual cells.

5.3 Post-processing Based Approach

In native XML DB, access control through post-processing described as:

$$SA_X = ACR\langle A_X \rangle = ACR\langle Q_X\langle D_X \rangle \rangle$$

In XRDB, this approach could be conducted through: (1) XML answer filtering (③ in Fig. 3); or (2) relational answer filtering (⑥ in Fig. 3). (1) is similar to the postprocessing approach in [55], while (2) evaluates relational query Q_R to obtain unsafe relational answer, and process ACR_R against the answers:

$$SA_X = C^{-1}(SA_R) = C^{-1}(ACR_R\langle A_R \rangle) = C^{-1}(ACR_R\langle Q_R\langle D_R \rangle \rangle)$$

However, the post-processing filters require the intermediate answers ($\langle A_R \rangle$ or $\langle A_X \rangle$) to retain information of the original paths for ACR to operate on. As an example of this approach, [28] check streaming XML data against both query and ACR at the same time. Since it works in the streaming data environment, full paths are retained. However, in most X2R conversion algorithms, the intermediate answer A_R or A_X does not contain full path information. Therefore, postprocessing approaches are not suitable for all applications.

6 Experimental Validation

To show that the proposed theory and implementations are practical yet efficient, we show our preliminary experimental results.

An XML document with 8517 nodes are generated by XMark [50]. We use XRDB(XRel) [48], with Oracle 10g as underlying RDBMS. We design five roles, abbreviated as A (administrator), M , RU , S and U , respectively. We do not have any descendant elimination negative rules since XRDB(XRel) cannot directly handle it (Lemma 3). We generate four groups of synthetic XPath queries, each has a different setting of wildcards and predicates.

In the XRDB(XRel), we convert all rules into relational, and enforce them through views and VPD. For a comparison, we also enforce same rule sets on the same XML document in native XML environment. We enforce XML access

Fig. 5. Query processing time for four sets of queries

control rules using QFilter [24], and answer XML queries using Galax. In all the experiments, we use the *query processing time* as an evaluation metric. Figure 5 shows the results. Comparing both view-based and VPD-based approaches with the reference (no security enforcement), our approaches do not add much overhead for fine-grained access control. Meanwhile, the size of accessible data gets smaller with security enforcement. Therefore, querying on smaller set of records is even faster than that on no-security case.

7 Conclusion

In this paper, we propose a generic analysis to the access control problem in XRDB. We first analyze XML control models to propose a formal description of XML access control using deep set operators. Then we articulate the problem of XML access control in XRDB as essentially the problem of XML/Relational object and operation equivalency and conversion. We show that, equivalent counterparts of deep set operators in relational model are needed to fully implement XML access control in XRDB. We analyze the definition and semantics of each operator, and show how they can be converted to XRDB through two lemmas. Although detailed conversion implementation is connected with the specific X2R conversion algorithm used in XRDB, we propose an algebraic description of these operators. Moreover, we study possible implementations of XML access control in XRDB. We categorize them into three approaches, and formally describe the semantics of each approach using deep set operators. Finally, we show the validity of our approaches using experiment results.

References

1. Bray, T., Paoli, J., Sperberg-McQueen (Eds), C.M.: Extensible Markup Language (XML) 1.0 (2nd Ed.). W3C Recommendation (2000)
2. Godik, S., Moses (Eds), T.: eXtensible Access Control Markup Language (XACML) Version 1.0 (2003)
3. Damiani, E., De Capitani di Vimercati, S., Paraboschi, S., Samarati, P.: A Fine-Grained Access Control System for XML Documents. *ACM TISSEC* **5**(2) (May 2002) 169–202
4. Bertino, E., Ferrari, E.: Secure and Selective Dissemination of XML Documents. *ACM TISSEC* **5**(3) (August 2002) 290–331
5. Tan, K.L., Lee, M.L., Wang, Y.: Access Control of XML Documents in Relational Database Systems. In: IC, Las Vegas, NV (June 2001)
6. Barbosa, D., Freire, J., Mendelzon, A.O.: Designing Information-preserving Mapping Schemes for XML. In: VLDB, Trondheim, Norway (2005) 109–120
7. Samarati, P., Bertino, E., Jajodia, S.: “An Authorization Model for a Distributed Hypertext System”. *IEEE TKDE* **8**(4) (1996) 555–562
8. Damiani, E., Vimercati, S.D.C.D., Paraboschi, S., Samarati, P.: Design and Implementation of an Access Control Processor for XML Documents. *Computer Networks* **33**(6) (2000) 59–75
9. Kudo, M., Hada, S.: XML Document Security Based on Provisional Authorization. In: ACM CCS. (2000)
10. Fundulaki, I., Marx, M.: Specifying access control policies for xml documents with xpath. In: ACM SACMAT. (2004) 61–69
11. Fernandez, E., Gudes, E., Song, H.: A Model of Evaluation and Administration of Security in Object-Oriented Databases. *IEEE TKDE* **6**(2) (1994) 275–292
12. Wang, J., Osborn, S.L.: A role-based approach to access control for XML databases. In: ACM SACMAT. (2004) 70–77
13. Bertino, E., Castano, S., Ferrari, E.: Securing XML Documents with Author-X. *IEEE Internet Computing* **5**(3) (2001) 21–31
14. Murata, M., Tozawa, A., Kudo, M., Hada, S.: XML access control using static analysis. *ACM TISSEC* **9**(3) (2006) 292–324
15. Damiani, E., di Vimercati, S.D.C., Paraboschi, S., Samarati, P.: Securing XML Documents. In: EDBT. (2000) 121–135
16. Cho, S., Amer-Yahia, S., Lakshmanan, L.V., Srivastava, D.: Optimizing the Secure Evaluation of Twig Queries. In: VLDB, Hong Kong, China (August 2002)
17. Xiao, Y., Luo, B., Lee, D.: Security-Conscious XML Indexing. In: DASFAA, Bangkok, Thailand (2007)
18. Yu, T., Srivastava, D., Lakshmanan, L.V., Jagadish, H.V.: Compressed Accessibility Map: Efficient Access Control for XML. In: VLDB, Hong Kong, China (2002)
19. Jiang, M., Fu, A.W.C.: Integration and Efficient Lookup of Compressed XML Accessibility Maps. *IEEE TKDE* **17**(7) (2005) 939–953
20. Stoica, A., Farkas, C.: Secure XML Views. In: DBSec. (2002) 133–146
21. Fan, W., Chan, C.Y., Garofalakis, M.: Secure XML querying with security views. In: SIGMOD. (2004) 587–598
22. Kuper, G., Massacci, F., Rassadko, N.: Generalized XML security views. In: SACMAT. (2005) 77–84
23. Murata, M., Tozawa, A., Kudo, M.: XML Access Control using Static Analysis. In: ACM CCS, Washington D.C. (2003)

24. Luo, B., Lee, D., Lee, W.C., Liu, P.: QFilter: Fine-Grained Run-Time XML Access Control via NFA-based Query Rewriting. In: ACM CIKM, Washington D.C., USA (November 2004)
25. Qi, N., Kudo, M.: Access-condition-table-driven access control for xml databases. In: ESORICS. (2004) 17–32
26. Qi, N., Kudo, M.: Xml access control with policy matching tree. In: ESORICS. (2005) 3–23
27. Mohan, S., Sengupta, A., Wu, Y.: Access control for XML: a dynamic query rewriting approach. In: CIKM. (2005) 251–252
28. Bouganim, L., Ngoc, F.D., Pucheral, P.: Client-Based Access Control Management for XML Documents. In: VLDB, Toronto, Canada (2004)
29. Bertino, E., Ferrari, E., Provenza, L.P.: Signature and Access Control Policies for XML Documents. In: ESORICS. (2003) 1–22
30. Carminati, B., Ferrari, E., Bertino, E.: Securing XML data in third-party distribution systems. In: CIKM. (2005) 99–106
31. Finance, B., Medjdoub, S., Pucheral, P.: The case for access control on XML relationships. In: CIKM. (2005) 107–114
32. Mohan, S., Wu, Y.: IPAC: an interactive approach to access control for semi-structured data. In: VLDB, VLDB Endowment (2006) 1147–1150
33. Jajodia, S., Sandhu, R.: Toward a Multilevel Secure Relational Data Model. In: SIGMOD. (May 1990)
34. Winslett, M., Smith, K., Qian, X.: Formal Query Languages for Secure Relational Databases. ACM TODS **19**(4) (1994) 626–662
35. Sandhu, R., Chen, F.: The Multilevel Relational (MLR) Data Model. ACM TISSEC **1**(1) (1998)
36. Griffiths, P.P., Wade, B.W.: An Authorization Mechanism for a Relational Database System. ACM TODS **1**(3) (September 1976) 242–255
37. Jajodia, S., Samarati, P., Sapino, M.L., Subrahmanian, V.S.: Flexible Support for Multiple Access Control Policies. ACM TODS **26**(2) (June 2001) 214–260
38. Jajodia, S., Samarati, P., Subrahmanian, V.S., Bertino, E.: A Unified Framework for Enforcing Multiple Access Control Policies. In: ACM SIGMOD. (May 1997) 474–485
39. Gabillon, A., Bruno, E.: Regulating access to XML documents. In: DBSec. (2002) 299–314
40. Murthy, R., Liu, Z.H., Krishnaprasad, M., Chandrasekar, S., Tran, A.T., Sedlar, E., Florescu, D., Kotsovolos, S., Agarwal, N., Arora, V., Krishnamurthy, V.: Towards an enterprise XML architecture. In: ACM SIGMOD. (2005) 953–957
41. Rys, M.: XML and relational database management systems: inside Microsoft SQL Server 2005. In: ACM SIGMOD. (2005) 958–962
42. Nicola, M., van der Linden, B.: Native XML support in DB2 universal database. In: VLDB. (2005) 1164–1174
43. Beyer, K., Ozcan, F., Saiprasad, S., der Linden, B.V.: DB2/XML: designing for evolution. In: ACM SIGMOD. (2005) 948–952
44. Deutsch, A., Fernandez, M.F., Suciu, D.: Storing Semistructured Data with STORED. In: ACM SIGMOD, Philadelphia, PA (Jun. 1998)
45. Shanmugasundaram, J., Tufte, K., He, G., Zhang, C., DeWitt, D., Naughton, J.: Relational Databases for Querying XML Documents: Limitations and Opportunities. In: VLDB, Edinburgh, Scotland (September 1999)
46. Lee, D., Chu, W.W.: Constraints-preserving Transformation from XML Document Type Definition to Relational Schema. In: ER, Salt Lake City, UT (2000) 323–338

47. Florescu, D., Kossmann, D.: Storing and Querying XML Data Using an RDBMS. "IEEE Data Eng. Bulletin **22**(3) (Sep. 1999) 27–34
48. Yoshikawa, M., Amagasa, T., Shimura, T., Uemura, S.: XRel: A Path-Based Approach to Storage and Retrieval of XML Documents using Relational Databases. ACM TOIT **1**(2) (November 2001) 110–141
49. Lee, D., Lee, W.C., Liu, P.: Supporting XML Security Models using Relational Databases: A Vision. In: XSym, Berlin, Germany (September 2003)
50. Schmidt, A.R., Waas, F., Kersten, M.L., Florescu, D., Manolescu, I., Carey, M.J., Busse, R.: The XML Benchmark Project. Technical Report INS-R0103, CWI (April 2001)
51. Qi, N., Kudo, M., Myllymaki, J., Pirahesh, H.: A function-based access control model for xml databases. In: ACM CIKM. (2005) 115–122
52. Luo, B., Lee, D., Lee, W.C., Liu, P.: Deep Set Operators for XQuery. In: ACM SIGMOD Workshop on XIME-P, Baltimore, MD, USA. (2005)
53. Berglund, A., Boag, S., Chamberlin, D., Fernandez, M.F., Kay, M., Robie, J., Simeon, J.: XML Path Language (XPath) 2.0. W3C Working Draft (November 2003)
54. Boag, S., Chamberlin, D., Fernandez, M.F., Florescu, D., Robie, J., Simeon, J.: XQuery 1.0: An XML Query Language. W3C Working Draft (November 2003)
55. Luo, B., Lee, D., Lee, W.C., Liu, P.: A Flexible Framework for Architecting XML Access Control Enforcement Mechanisms. In: VLDB Workshop on SDM, Toronto, Canada (2004)
56. Luo, B., Lee, D., Liu, P.: Pragmatic XML access control using off-the-shelf RDBMS. Technical report, Penn State University (2007)
57. Lu et al., H.: What makes the differences: benchmarking XML database implementations. ACM TOIT **5**(1) (2005) 154–194