

# Incremental Adaptation of XPath Access Control Views

Padmapriya Ayyagari, Prasenjit Mitra, Dongwon Lee, Peng Liu<sup>\*</sup>, Wang-Chien Lee<sup>†</sup>  
The Pennsylvania State University  
University Park, PA, USA 16802.

{payyagari,pmitra,dlee,pliu}@ist.psu.edu,wlee@ist.psu.edu

## ABSTRACT

Materialized XPath access-control views are commonly used for enforcing access control. When access control rules defining a materialized XML access-control view change, the view must be adapted to reflect these changes. The process of updating a materialized view after its definition changes is referred to as *view adaptation*. While XPath security views have been widely reported in literature, the problem of view adaptation for XPath security views has not been addressed. View adaptation results in view downtime during which users are denied access to security views to prevent unauthorized access. Thus, efficient view adaptation is important for making XPath security views pragmatic. In this work, we show how to adapt an XPath access-control view incrementally by re-using the existing view, which reduces computation and communication costs significantly, and results in less downtime for the end-user. Empirical evaluations confirm that the incremental view adaptation algorithms presented in this paper are efficient and scalable.

## Categories and Subject Descriptors

H.2.7 [Database Management]: Database Administration Security, integrity, and protection; K.6.5 [Management of Computing and Information Systems]: Security and Protection

## General Terms

Security

## Keywords

XML Access Control, XPath View, View Adaptation

<sup>\*</sup>Peng Liu was supported in part by NSF CCR-TC-0233324 and NSF/DHS 0335241

<sup>†</sup>Wang-Chien Lee was supported in part by National Science Foundation grants IIS-0328881, IIS-0534343 and CNS-0626709

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ASIACCS'07, March 20-22, 2007, Singapore.

Copyright 2007 ACM 1-59593-574-6/07/0003 ...\$5.00.

## 1. INTRODUCTION

Access control is an important component of maintaining the security of data and information systems. Access control rules [6, 15, 19, 8] specify the parts of the data that users can access and parts that they cannot. Materialized access-control views are an efficient way of implementing access control. A materialized access-control view contains only the data to which an user or role has access. User queries are evaluated against this access-control view without the need for any further security check. In this paper, we use the terms “access-control views” and “security views” interchangeably.

Increasingly, XML is being used as a format for storing and exchanging data in a large number of applications such as web-based applications, e-commerce, and XML client-server databases. In all these applications, there is a need for access control. Recently, there is also an increasing interest in supporting fine-grained XML access control [6]. Towards this end, techniques such as XPath security views [7, 13, 25] are being used. Client-side XPath security views [5] and semantic caching [28] are popular in decentralised systems. Such remote views reduce communication costs by allowing user queries to be evaluated locally, but add to communication costs incurred while maintaining the view.

Two issues arise in the case of materialized views - one is the **view update** problem [11] or the problem of synchronizing materialized views with updates to the *base data*, and the other is the **view adaptation** problem [12] or the problem of synchronizing materialized views with changes to the view definition. This paper deals with the latter. Due to changes in organizational policies or security requirements, access control rules for a user (or role) can be updated from time to time. When access control rules change, the corresponding access-control view definition also changes, and the view must be *adapted* [12] to conform to the user's new access permissions. When access-control views are being adapted, users are denied access to the view to ensure that no unauthorized access to data occurs. This “downtime” could have important, often commercial, implications for several applications.

XPath access-control views have been widely reported in literature; however the problem of adapting them efficiently has not been addressed. Recently, work has been reported on updating XML views [14]. However, to the best of our knowledge, no work has been done on the problem of adapting XPath views. We believe that to make materialized security views pragmatic, efficient adaptation of such views is vital. This is because, in numerous applications and environments, there might be frequent updates to access control rules neces-

sitating view adaptation. We outline two such scenarios below to motivate the *efficient XPath view adaptation* problem.

**Example 1 (Dynamic Environment).** In dynamic environments such as distributed, heterogeneous and web-based environments that use XML to store and integrate information from multiple sources, changes in the access rules and policies of the underlying data sources are not uncommon [26]. This could be the result of the dynamic nature of various organizational policies, changing roles, security requirements, etc. Such changes could necessitate changes to access-control views defined on the data. □

**Example 2 (Schema Evolution).** In applications like GIS, schemas evolve frequently [21], thus the semantics of the XPath views also change, necessitating view adaptation (discussed in Section ??). Mergers of organizations, establishment of new business alliances, changes in organizational structure, etc. are examples of scenarios where schema definitions, access control rules and policies undergo several changes before things stabilize [10]. In such cases, views would need to be adapted frequently. □

In both of the scenarios, changes to view definitions are frequent enough to result in downtime that is a nuisance to the end-user. However, there are substantial performance gains from materializing the views. In order to keep the downtime to a minimum, efficient view adaptation is necessary. There are two factors to be considered while adapting views efficiently: (1) *Computation cost* for adapting access-control views, which is critical for interactive applications [12]; and (2) *Communication cost* for transferring data to adapt remote security views, e.g., in client-server systems, and client-based access control [5].

When view definitions change, obvious but naive way is to simply re-compute the views from the base data. An alternative is to “incrementally” adapt the view by fetching only the data that is not already present in the view. We refer to this technique as **incremental view adaptation**. [12]. That is, the incremental adaptation of XPath access-control views is the focus of our work. Incremental view adaptation reduces computation and communication costs by exploiting the following observations:

**Observation 1** *If a new view is entirely contained in an old view, no changes to the materialized view are necessary.* □

**Observation 2** *If a new view is partially overlapping with an old view, only the difference needs to be materialized.* □

Considerable work has been done on the view update problem in the relational domain [24]. Gupta et al. [12] have reported work on incrementally adapting relational views. Due to the semi-structured and hierarchical nature of XML, the problem of adapting XPath views is different from the problem of view adaptation in relational databases. Furthermore, XPath access-control views present additional challenges because typically these views are defined using both positive and negative rules [6, 15, 8]. Typically, non-security views do not have negative rules in their definitions. The presence of negative access control rules makes adaptation more complicated for access-control views. To the best of our knowledge, no prior work has addressed the problem of XPath view adaptation.

Our key contributions in this paper are as follows:

1. We present a set of comprehensive incremental view adaptation techniques for XPath access-control views expressed for the  $XP\{/,//,*,\square\}$  (the fragment of XPath with step, descendant, wild-card and predicate operators, defined in [18]) that reduce computation cost and optimize communication cost. We also exploit XPath containment [17] and query answering using views [28] for efficient adaptation; these issues have not been addressed by previous work in relational databases [12].
2. We also suggest auxiliary information that can be stored to improve the performance of our incremental view adaptation techniques.
3. We demonstrate the superiority of incremental view adaptation over view re-computation experimentally.

The rest of the paper is organized as follows. In Section 2, we discuss related work. We present background information and details about the setting where our algorithms can be applied in Section ?? . We define the XPath view adaptation problem in Section 4. The incremental view adaptation algorithms are presented in Section 5. We present a validation of our algorithms in Section 7. The paper is concluded in Section 8.

## 2. RELATED WORK

XPath views for implementing security or access-control for XML documents has been widely reported in literature. Stoica and Farcas first proposed XPath security views [25]. Fan et al. [7] propose algorithms for deriving security view definitions from security policies, and present techniques for efficient query processing on security views. Kuper et al. [13] further generalize the notion of XML security views by creating a “view DTD” that hides the base data DTD from users. Yu et al. propose a method using compressed XML views to support access control [29] for reducing the storage overhead for materialized views.

Bouganim et al. [5] talk about client-based access control for XML documents where the access-control is monitored at the client side. This model is suitable for environments where clients do not place sufficient trust in their data service provider.

Several models for specifying access control rules for XML access control have been mentioned. The model proposed by Damiani et al. [6] is very popular and Fundulaki et al. [8], Fan et al. [7], Kuper et al. [13] and Luo et al. [15] propose models that are similar to [6]. Lim et al. [14] propose a more sophisticated priority-based access control model. Recently, some new standards for XML access control, such as XACML have also emerged.

Gupta, et al., [12] originally introduced the problem of view adaptation for relational databases. They consider all possible “redefinitions” of SQL SELECT-FROM-WHERE-GROUPBY-HAVING, UNION, EXCEPT views and show how such views can be adapted using old materializations wherever possible. They identify auxiliary information to be stored with the views to assist redefinitions. Bellahsene also proposes a “fragment-based” approach for relational databases to view adaptation in a setting where fragments of a view are materialized instead of the complete view [28]. They propose a method where adaptation is performed using not only the old view but all the materialised views in the system.

A significant amount of work has been done in the area of XML view maintenance problem when the data in the database changes [22, 14]; however, to the best of our knowledge, no work has been done on the view adaptation problem for XPath views.

Chen, et al., [38] have proposed a cache-aware XQuery answering system that checks whether a query can be answered using cached results of previous queries. This problem is similar to ours in that the cached query results are similar to our views. Their algorithm is based on the existence of a single containment mapping from the query to the view and a set of heuristics. Consequently, their algorithm is not complete. Balmin, et al., [4] have also presented sound but incomplete algorithms for using materialized XPath views for query processing using XML values, full paths and node references. Mandhani and Suciu [17] and Xu and Ozsoyoglu [28] have shown (independently) how a query can be answered using a view. We use their technique in our view adaptation algorithm. Note that there exists no algorithm in the literature that can compute the rewriting of an XPath query using multiple views. However, when an algorithm to rewrite an XPath query using multiple views is developed, our algorithms can utilize it seamlessly.

Miklau and Suciu addressed the XPath containment problem and prove that it is co-NP-complete [18]. They provide a sound and complete exponential time algorithm for XPath containment, and some parameterized polynomial time algorithms. Neven and Schwentick [20] have shown the complexity of XPath containment in the presence of disjunctions and DTDs. Schwentick [23] provides algorithms to test for the containment of XPath in the presence of disjunctions using several techniques, like tree automata.

Note that these complexities are typically with respect to the size of the XPath query and the size of the view definitions, *and not in the size of the data in the view or the documents in the database*. Hence, in practice, the algorithms utilizing these techniques complete in reasonable time.

### 3. XPATH VIEW-BASED ACCESS CONTROL

To make our view adaptation technique pragmatic, we adopt a popular, very commonly used XML-access-control language similar to the models proposed by [6, 15, 8, 19] for specifying access control rules.

An XML document is represented as a hierarchy of nested nodes (elements and attributes) and fine-grained access control is established at the node level. In our model, the node-level authorization is specified via 5-tuple *access control rules* [6, 15, 8]:

*access control rule* := {*subject, object, action, sign, type*}

where (1) *subject* is to whom an access is granted or denied (depending upon the sign of the rule: + or -), i.e., user or role; (2) *object* refers to nodes in XML documents specified by an XPath expression (XPath can be used to identify nodes in an XML document); (3) *action* is one of “read”, “write”, and “update”; (4) *sign*, +, - indicates whether access is “granted” (positive rule) or “denied” (negative rule), respectively; and (5) *type* LC, RC refers to either Local Check (i.e., access is granted or denied to only attributes or textual data of nodes in context, i.e., `self::text() | self::attribute()`

in XPath), or Recursive Check (i.e., access is granted or denied to current nodes and propagated to all their descendants `descendant-or-self::node()`, respectively). For a node in a document to be in a view, under the recursive check semantics, an ancestor of the node must be included by a positive rule and none of its ancestors should be excluded by a negative rule.

By default, access is denied to all nodes whose authorizations are not specified, either explicitly (via LC rules) or implicitly (via RC rules). A node can have more than one relevant rule. If a conflict occurs between a positive (+) and negative (-) rule, the negative rule takes precedence. For our work, we consider only *read-action rules with recursive check*. Consequently, in the rest of the paper, we will refer to only the object and the sign of a rule.

**Example 3.** Consider the view *V* shown in Figure 1(c) defined on *D* from Figure 1(a). If the view is defined using only the positive rule: */Employee/EmployeeRecord* and a negative rule */Employee/EmployeeRecord/Pay*, then *V* would contain *EmployeeRecord* and all its children, but not the *Pay* element, or its *Basic* and *Allowance* children. □

Changes to access control rules are not the only reason why views may need to be adapted. To explain how changes to XML schema induce changes in view definitions, consider the following example. Suppose the IDs for all the employees were updated to include their social security number, then views defined using the positive rule */Employee/EmployeeRecord/profile/\** would have access to social security information which was not originally intended. If (some of) the views are not supposed to have access to such confidential information, the system administrator will have to add */site/people/person/profile/ID* as a negative rule for all those views (that do not already have this as a negative rule). This would induce changes in view definitions and would necessitate view adaptation. Such changes would also be common in applications like GIS where schemas are dynamic (dynamic schema evolution) [21].

### 4. PROBLEM DEFINITION

An access-control view contains parts of XML documents that are included by the positive (+) access control rules (denoted  $ACR^+$  but **not** included by the negative (-) access control rules (denoted  $ACR^-$ ).

Given a set of positive access control rules  $ACR^+ = p_1, p_2, \dots, p_n$ , and set of negative access control rules,  $ACR^- = n_1, n_2, \dots, n_k$ , the corresponding access control view *V* is defined by the expression:

$$V = (p_1(D) \cup p_2(D) \dots \cup p_k(D)) -^D (n_1(D) \cup n_2(D) \dots \cup n_k(D))$$

where  $p_i$  and  $n_i$  are XPath expressions belonging to the fragment of XPath:  $XP^{\{/, //, *, []\}}$  and  $-^D$  is the deep-except operator (defined in Section 4.2). Updates that cause view definition change include the following events: (1) *Removal of a positive rule*: removing that the view previously had access to. (2) *Addition of a positive rule*: potentially adding data. (3) *Removal of a negative rule*: potentially allowing access to more data. (4) *Addition of a negative rule*: potentially removing access to data. Formally, we study the following problem:

**Definition 1 (XPath access control view adaptation problem)**  
Given an access control view *V*, whose definition changes due

```

D:
<Employee>
  <EmployeeRecord category = "...">
    <Profile>
      <Name> ... </Name>
      <Age> ... </Age>
      <Pay>
        <Basic> ... </Basic>
        <Allowance> ... </Allowance>
      </Pay>
      <Department> ... </Department>
      <ID> ... </ID>
    </Profile>
  </EmployeeRecord>
</Employee>

```

(a)

```

V:
<Employee>
  <EmployeeRecord>
    <Profile>
      <Name> ... </Name>
      <Age> ... </Age>
    </Profile>
  </EmployeeRecord>
</Employee>

```

(b)

```

V:
<Employee>
  <EmployeeRecord category = "...">
    <Profile>
      <Name> ... </Name>
      <Age> ... </Age>
      <Department> ... </Department>
      <ID> ... </ID>
    </Profile>
  </EmployeeRecord>
</Employee>

```

(c)

Figure 1: A running example showing: (a) XML document (b) XML security view. (c) Recursive Check.

to an update of one of the four types specified above to the corresponding access control rules, adapt the view  $V$  such that it satisfies the confidentiality criteria (defined in Section 4.1).

## 4.1 Security Criteria

We have two main security criteria:

- **Confidentiality:** At time  $t$ , the system's access-control state consists of two important elements: (C1) the current set of access-control rules including all the changes that have been proposed by time  $t$ , and, (C2) the set of views that are currently accessible. Note that C2 does not include views that are NOT accessible at time  $t$  (because they are undergoing adaptation).

The confidentiality criteria says that at any point of time  $t$ , C2 will not violate C1, i.e., no accessible view in C2 will contain data the user is not allowed to access at time  $t$  by C1.

- **Denial of service (or downtime):** When access control rules defining a view undergo a change threatening confidentiality, the view is usually "shut down" to prevent loss of confidentiality, adapted, and then put back online. During this time, users are denied access to the view.

## 4.2 XPath Set and Deep-Set Operators

In this section, we outline operators that are used for manipulating XPath documents.

For the XPath *union* operator, we follow the same semantics as defined by the W3C [3]. That is, the semantics of  $X1(D) \cup X2(D)$  is defined as the union of two node sequences returned by  $X1(D)$  and  $X2(D)$ .  $X1(D) \cup X2(D)$  takes two node sequences as operands and returns a sequence containing all the nodes that occur in either of the operands [3]. The standard XPath *intersect* operator takes two node sequences as operands and returns a sequence containing all the nodes that occur in both operands [3].

However, in the context of XML access control with recursive semantics where the views are defined using not only




Figure 2: Example XML data  $D$  and two rules  $/x/a[e]$  and  $//d[f]$

the nodes but the entire subtrees rooted at the nodes, the formal semantics of regular XPath set operators is not sufficient. Suppose, as shown in Figure 2, there is a positive rule  $R1 : (admin, /x/a[e], read, +, RC)$  and a negative rule  $R2 : (admin, //d[f], read, -, RC)$ . Then, the *admin* can read the data specified by  $/x/a[e]$  'minus' the data specified by  $//d[f]$ . However, the default except operator of XPath cannot capture the notion of 'minus' correctly because it is based on the node-IDs and ignores subtrees completely. Here, data corresponding to rules  $R1$  and  $R2$  are indicated by the two rectangular boxes. The *admin* should have access only the nodes with node-IDs: 2, 3, 4, 8. However, when the 'minus' is expressed using the XPath except operator,  $/x/a[e] \text{ except } //d[f]$  is operated in terms of node-IDs as in  $\{2\} \text{ except } \{5\} = \{2\}$ . Then, the returned answers are the nodes with node-ID with '2' and all of its descendants 3, 4, 5, 6, 7, 8. However, this is problematic because the answer contains the nodes 5, 6, 7 violating access control rules.

Similarly, we need to extend the *intersect* operator as well. The deep-set operators with extended semantics are denoted as deep-except ( $-^D$ ), and deep-intersect ( $\cap^D$ ). For




Figure 3: eXPath set and deep-set operators.

instance, the correct semantics for  $P1 \stackrel{-D}{-} P2$  is: (1) when answers to  $P2$  are descendants of those to  $P1$ , nodes in the answer to  $P2$  are excluded from the answer to  $P1 \stackrel{-D}{-} P2$ ; and (2) otherwise, it degenerates to the regular *except* operator,  $P1 \text{ except } P2$ . As an example, Figure 3 illustrates different results of the same query when both standard and extended set operators are used for  $D$  of Figure 2. Hereafter, we assume that there exist efficient implementations of deep-set operators. In our experimentations, we used the implementation provided by Luo, et al.,[16]. Since they are implemented as user-defined functions of XML databases, they do not require any extra support from underlying XML databases.

### 4.3 XPath Containment and Rewriting

When a view definition is updated, our algorithm must determine whether the new view can be constructed using the existing views. If the view can be adapted using the old views, then there is no need to fetch data from the database and the incremental adaptation can be performed quickly. The problem of identifying which nodes in the old XPath views satisfy the new rule (added or deleted) is closely related to the problem of answering XPath queries using XPath views [28], which, in turn is closely related to the problem of XPath containment [20]. We define these problems below.

We have adapted the definition of XPath containment provided by Neven and Schwentick [20]:

**Definition 2 (XPath Containment)**  $p$  is contained in  $q$ , if for all XML documents, for all nodes  $t \in p$ , then  $t \in q$ . That is, if  $p$  is contained in  $q$ , then, if  $t$  is an answer to  $p$ , then  $t$  is also an answer to  $q$ .  $\square$

**Example 4.** Given the following queries:  $p = \text{/Employee/EmployeeRecord/Pay/Basic}$  and  $q = \text{/Employee/EmployeeRecord/Pay/*}$ ,  $p$  is contained in  $q$ . For  $q = \text{/Employee/EmployeeRecord/age > 50/Pay/Basic}$ ,  $p$  is not contained in  $q$ .  $\square$

We define the query rewriting problem as follows:

**Definition 3 (Query Rewritability Using a View)** Given a query pattern  $q$  and a view  $v$ , does there exist a compensation pattern  $q'$  such that for every database  $D$  and view extension  $\mathcal{E}$  of  $v$ ,  $q'(\mathcal{E}) \subseteq q(D)$ ?  $\square$

**Example 5.** Let  $D$  be the XML Document  $D$  in Figure 1,  $V = \text{/Employee/EmployeeRecord/Pay}$ ,  $Q = \text{/Employee/EmployeeRecord/Pay[/Basic]}$ , and  $C = \text{Pay[/Basic]}$ . Executing the *compensation query*  $C$  on  $V(D)$  yields results that are also obtained by executing  $Q$  on  $D$ .  $\square$

Suppose  $Q$  above was a negative rule being added to the view definition  $V$  currently defined by a single positive rule

$\text{/Employee/EmployeeRecord/profile/Pay}$ .  $Q$  is re-writable using  $V(D)$  and the compensation query  $C$ . So  $Q$  can be executed on  $V(D)$  to determine what parts of  $V(D)$  must now be removed due to the addition of the negative rule  $Q$ .

Note that although we know that when a negative rule is added, we are removing data from the view that is already contained in the view, it might not always be possible to determine which nodes constitute that data. Consider the following example.

**Example 6.** Let  $ACR^+ = \{\text{/Employee/EmployeeRecord/profile/Pay}\}$ ,  $ACR^- = \{\}$ , and,  $R = \text{/Employee/EmployeeRecord/profile[age > 50]/Pay}$ . Although we know that the  $\langle \text{Pay} \rangle$  elements for all employees (including the ones aged above 50) are contained in the view, when the negative rule  $R$  is added, it is not possible to determine from the view alone which  $\langle \text{Pay} \rangle$  elements must be deleted and which must remain in the view. This is because  $R$  is not rewritable using the view, although it is contained in the view!  $\square$

**Example 7.** Consider a view that is defined using the following rules:  $ACR^+ = \{\text{/Employee/EmployeeRecord/profile/Pay, /Employee/EmployeeRecord/profile/age}\}$ , and  $ACR^- = \{\}$ .  $R = \text{/Employee/EmployeeRecord/profile[age > 50]/Pay}$  is a negative rule being added to the view. Then using the positive rule  $\text{/Employee/EmployeeRecord/profile/Pay}$  or  $\text{/Employee/EmployeeRecord/profile/age}$  alone,  $R$  is not rewritable. To rewrite  $R$  using multiple views, one possible approach is to store node IDs for each  $\langle \text{Pay} \rangle$  element and  $\langle \text{age} \rangle$  element to determine their common  $\langle \text{profile} \rangle$  parent and store them together under  $\langle \text{profile} \rangle$  element (as shown in Example 9 below). This is because, for each  $\langle \text{Pay} \rangle$  element the corresponding  $\langle \text{age} \rangle$  element can be determined and checked to see if its value is greater than 50. If it is, then the corresponding  $\langle \text{Pay} \rangle$  element would appear in the result set. The path from the root for the profile element must be stored as well to determine the parent of each element. The view contains:

```

<profile>(path : /Employee/EmployeeRecord/profile)
  <Pay>... </Pay>
  <age>... </age>
</profile>
<profile>(path : /Employee/EmployeeRecord/profile)
  <Pay>... </Pay>
  <age>... </age>
</profile>

```

If the view stores all the profile elements with their  $\langle \text{Pay} \rangle$  and  $\langle \text{age} \rangle$  children as shown above, then the compensation query for  $R$  would be  $\text{/profile[age > 50]/Pay}$ .  $\square$

However, sometimes, although an expression is re-writable using the positive rules in a view, the presence of negative rules might still make not allow the expression to be rewritten using the view.

**Example 8.** Consider the following example where a negative rule  $R$  is added to the view:  $ACR^+ = \text{/Employee/EmployeeRecord/profile/*}$ ,  $ACR^- = \text{/Employee/EmployeeRecord/profile/age}$ .  $R = \text{/Employee/EmployeeRecord/profile[age > 50]/name}$  is a negative rule being added to the view.  $R$  is rewritable using  $\text{/Employee/EmployeeRecord/profile/*}$  ( $ACR^+$ ). However, due to  $\text{/Employee/EmployeeRecord/profile/age}$  ( $ACR^-$ ), the view does not contain the  $\langle \text{age} \rangle$  element of the employee's profile. So if  $R$ , which is re-writable using  $ACR^+$ , is executed

on the view, it would not be able to determine which name elements belong to employees whose  $\langle age \rangle$  is above 50, and which elements belong to employees who belong to the IST department. Hence R is contained in the view, is re-writable using  $ACR^+$ , but is not answerable using the view.  $\square$

**Example 9.** Consider the view defined using the following rules:  $ACR^+ = \{/Employee/EmployeeRecord/profile/*\}$ , and,  $ACR^- = \{/Employee/EmployeeRecord/profile[age > 50]/pay\}$ , and, the query:  $R = \{/Employee/EmployeeRecord/profile[age < 25]/name\}$ . R is answerable using the view, because R is rewritable using  $ACR^+$  and R is not contained in  $ACR^-$ . However, if R was  $\{/Employee/EmployeeRecord/profile[age > 75]/name\}$  then R would not be answerable using the view.  $\square$

Since access control view definition can contain multiple positive and negative rules, XPath query containment and rewriting must be computed for unions of XPath expressions. Neven and Schwentick have shown that by adding disjunction the problem of containment of  $XP\{/,//,[],*\}$  remains in CO-NP [20]. Our algorithms (described in Section 5) use containment of  $XP\{/,//,[],*,*\}$  in the presence of unions. This containment can be checked using tree automata as discussed in [23].

The techniques proposed by Xu and Ozsoyoglu [28], and Mandhani and Suciu [17], can be used to obtain an equivalent rewriting of an XPath query using a single view. In their work, a view is defined using only one XPath expression. In the case of access control views, a view is defined using multiple positive (and negative) rules. In the simplest case, where no negative rules are present, an access control view, in our work, can be defined using multiple positive rules, which is equivalent to a *union of multiple views*.


#### 4.4 Cost Models

A view adaptation algorithm has the following costs: (a) *Computation cost*: The time taken to adapt the view. (b) *Communication cost*: The time taken to send the data needed for adaptation from the source database to a remote view.

In applications, where the materialized views are stored at the site of the base data itself, there is no or very little communication cost involved, the computation cost for view adaptation is important.

In systems, such as distributed or peer-to-peer databases, implementing client-based access control [5], views reside at the clients while the source database resides on a server (or remote site). Client-based access control is popular in (1) systems where clients do not trust database service providers to preserve data confidentiality, protect children from suspicious internet content, etc.; (2) decentralized data sharing systems like peer-to-peer databases. In such cases, the cost for actually computing the view might be insignificant when compared to the communication costs involved. In this case, there are fixed communication costs like latency, etc., and variable communication costs proportional to the size of the data. For large datasets, the latter costs dominate. Our algorithms do not have any effect on the latency of the communication anyway, so we focus on minimizing the data communicated to reduce the communication cost.

Although, the amount of data fetched from the database has an impact on the computation cost, in some cases compromising on the amount of data fetched from the database might lead to improved performance due to fewer operations for computing the adapted view.



**Figure 4: View Adaptation: (a) Adding a Positive rule R to  $ACR^+$ ; (b) Deleting a Negative rule R from  $ACR^-$ ; (c) Adding a Negative rule R to  $ACR^-$ ; and (d) Delete a Positive Rule R from  $ACR^+$ .**

## 5. XPATH VIEW ADAPTATION ALGORITHMS

In this section, we discuss the various view adaptation algorithms. We denote the base data as D.

### 5.1 Naive View Adaptation

In the naive method, the view is re-computed by executing the XPath expression in Equation ?? on the base data. We call this approach the naive view re-materialization approach. Although simple, this approach is not efficient as it involves many redundant computations (explained further in Section 5.2). The equation for the naive re-computation is given below:

$$(p_1(D) \cup p_2(D) \cup \dots \cup p_k(D)) -^D (n_1(D) \cup n_2(D) \cup \dots \cup n_k(D)) \quad (1)$$

where  $p_i \in ACR^+$  and  $n_i \in ACR^-$ .

### 5.2 Optimized View Adaptation

The expression for optimized view re-materialization is shown in Equation ?? below. This method is more efficient than the naive view re-materialization, because instead of computing all the XML nodes contained in  $(p_1(D) \cup p_2(D) \cup \dots \cup p_k(D))$  and  $(n_1(D) \cup n_2(D) \cup \dots \cup n_k(D))$  first, and then performing a deep-except operation on them, this scheme ignores all the (parts of) rules in  $ACR^-$  that do not intersect with  $ACR^+$  at all, and hence have no impact on the view. Although some negative rules might not intersect with any positive rules in the view definition, they are still defined to ensure that at a later point, when a positive rule is added, it does not bring in data that is not meant to be viewed by the user (role). The optimized equation for computing the contents of the adapted view is as follows:

$$(p_1(D) \cup p_2(D) \cup p_k(D)) -^D (ACR^- -^D p_i(D)) \quad (2)$$

where  $p_i \in ACR^+$

### 5.3 Incremental View Adaptation Under the Communication Cost Minimization Model

In this section, we discuss techniques for optimizing the communication cost for XPath security view adaptation. The examples used in this section are based on the schema in Figure 1. Figure 4 explains how we determine the incremental data to be added or removed from the view for all four cases.

In the equations provided below for the different cases, the ‘‘Default Rule’’ is executed at the database. Hence we have

provided the equations that minimize the data brought in from the database for the “Default Rules”. In all the other rules, we have stated the equations that minimize the computation.

We use the symbol  $ACR^{+new}$  to denote the new set of positive rules (after the addition or deletion of positive rules) and  $ACR^{-new}$  to denote the set of negative rules (after the addition or deletion of negative rules).

### 5.3.1 Addition of Positive Rules

Addition of a new positive rule R to  $ACR^+$  is processed using the following rules in order:

- **[Containment Rule]:** If R is contained in  $ACR^+ \cup ACR^-$ , i.e. R is either completely contained in a positive rule in the view, or is disallowed by some negative rule from the view, then do nothing and return.

**Example 10.** Let  $ACR^+ = \{/Employee/EmployeeRecord/profile/*$ ,  $ACR^- = \{/Employee/EmployeeRecord/profile/age$ , and  $R = \{/Employee/EmployeeRecord/profile/name$ . Then, R is already contained in  $ACR^+$ , and thus no changes to the view is necessary. Now, let  $R = \{/Employee/EmployeeRecord/profile/age$ . R is contained in  $ACR^-$ . No changes to the view are necessary because  $ACR^-$  prohibits R. □

If the containment rule is not applicable, then R potentially adds data to the view. Execute the default rule.

- **[Default Rule]:** Execute the following equation:

$$V'(D) = V(D) \cup (R(D) -^D (R'(D) \cup R''(D))) \quad (3)$$

where  $R' = R \cap^D ACR^+$  and  $R'' = R \cap^D ACR^-$

**Example 11.** Let  $ACR^+ = \{/Employee/EmployeeRecord/profile/name$ ,  $ACR^- = \{/Employee/EmployeeRecord/profile/age$ , and,  $R = \{/Employee/EmployeeRecord/profile/*$ . Then  $R' = \{/Employee/EmployeeRecord/profile/name$ , and  $R'' = \{/Employee/EmployeeRecord/profile/age$  in Equation ???. Upon execution of Equation ?? from the database, the expression adds the children of the node  $\{/Employee/EmployeeRecord/profile$  except the nodes  $\langle name \rangle$  that is already there and  $\langle age \rangle$  that is excluded by the negative rule. □

### 5.3.2 Deletion of Positive Rules

Deletion of a positive rule R from  $ACR^+$  is handled using the following rules in order:

- **[Containment Rule]:** If R is contained in  $ACR^{+new} \cup ACR^-$ , which means that removal of R from  $ACR^+$  does not remove any data from the view, do nothing and return.

**Example 12.** Let  $ACR^+ = \{/Employee/EmployeeRecord/profile/*$ ,  $\{/Employee/EmployeeRecord/profile/age$ ,  $ACR^{+new} = \{/Employee/EmployeeRecord/profile/*$ , and  $R = \{/Employee/EmployeeRecord/profile/age$ . R contained in  $ACR^{+new}$ , and  $ACR^{+new}$  is still in the view definition, so the view needs no change. □

**Example 13.** Consider the view defined as follows:  $ACR^+ = \{/Employee/EmployeeRecord/Profile/Name, \/Student/Profile/Name$  and  $ACR^- = \{\}$  (for simplicity). The view contains the names of employees and students of, say, an university. The  $\langle name \rangle$ -subtrees have ancestors:  $\langle \/Employee/EmployeeRecord/Profile \rangle$  or  $\langle \/Student/$

$Profile \rangle$ . When the rule  $R = \{/Student/Profile/Name$  is deleted, we know that all the answers to the query R, i.e., all student names are in the view but do not know which  $\langle name \rangle$ -subtrees have an  $\langle Employee/EmployeeRecord/Profile \rangle$  ancestor and which have a  $\langle Student/Profile \rangle$  ancestor and without additional information, there is no way to distinguish the two. Consequently, the rule has to be evaluated in the database using the Default Rule stated below. □

- **[Default Rule]:** The following equation minimizes the data communicated from the database to the client:

$$V'(D) = V(D) -^D (R(D) -^D (R'(D) \cup R''(D))) \quad (4)$$

where  $R' = R \cap^D ACR^{+new}$  and  $R'' = R \cap^D ACR^{-new}$ .

### 5.3.3 Addition of Negative Rules

Addition of a new negative rule R to  $ACR^-$  is processed using the following rules in order:

- **[Intersection Rule]:** If  $R \cap^D ACR^+ = \phi$ , then do nothing and return. R does not intersect with any positive rule in  $ACR^+$  and hence the addition of R does not affect the view.

**Example 14.** Let  $ACR^+ = \{/Employee/EmployeeRecord/profile/name$ ,  $ACR^- = \{/Employee/EmployeeRecord/profile/gender$ , and,  $R = \{/Employee/EmployeeRecord/profile/age$ . R does not intersect with any rule in  $ACR^+$ , so the view needs no update. □

- **[Containment Rule]:** If R is contained in  $ACR^-$ , then do nothing and return. R is contained in  $ACR^-$ . Thus, the view needs no change.

**Example 15.** Let  $ACR^+ = \{/Employee/EmployeeRecord/profile/name$ ,  $ACR^- = \{/Employee/EmployeeRecord/profile/gender, \/Employee/EmployeeRecord/profile/Pay/*$ , and  $R = \{/Employee/EmployeeRecord/profile/Pay/Allowance$ . R is contained in  $\{/Employee/EmployeeRecord/profile/Pay/$ . So the view needs no change. In the absence of predicates, the negative rule R can be simply executed on the view and the data satisfied by R can be deleted from the view. □

**Example 16.** Consider the view defined by the following positive rules:  $ACR^+ = \{/a/b, \/c/b$ . The view contains  $\langle b \rangle$ -subtrees whose parents are  $\langle a \rangle$  or  $\langle c \rangle$ . When the negative rule  $R = \{/a/b/e$  is added, we need to delete the  $\langle e \rangle$ -children of  $\langle b \rangle$ -children of  $\langle a \rangle$ -nodes. The view contains all the  $\langle b \rangle$ -children of  $\langle a \rangle$ -nodes and all the descendants of those  $\langle b \rangle$ -children. However, we do not know which  $\langle b \rangle$ -subtrees in the view have an  $\langle a \rangle$ -parent and which have a  $\langle c \rangle$ -parent. Consequently, we do not know which  $\langle e \rangle$ -nodes have to be removed from the view. Thus, the Default Rule has to be evaluated in the database server. □

- **[Default Rule]:** The default rule in this case is the same as the case where a positive rule was deleted: Equation ??.

### 5.3.4 Deletion of Negative Rules

Deletion of a negative rule  $R$  from  $ACR^-$  is processed using the following rules:

- **[Intersection Rule]:** If  $R \cap^D ACR^+ = \phi$ , do nothing and return.  $R$  does not intersect with any of the positive rules currently in  $ACR^+$  and hence removal of  $R$  does not affect the view.

**Example 17.** Let  $ACR^+ = \{/Employee/EmployeeRecord/profile/name\}$ ,  $ACR^- = \{\}$ , and, the negative rule being deleted is:  $R = \{/Employee/EmployeeRecord/profile/gender\}$ . Because the view does not have any gender information, the negative rule does not require the view to be adapted.

- **[Containment Rule]:** If  $R$  is contained in  $ACR^{-new}$ , do nothing and return. The removal of  $R$  does not affect the view because there is another negative rule that rules out all the data that  $R$  was ruling out.

**Example 18.** Let  $ACR^+ = \{/Employee/EmployeeRecord/profile/*\}$ ,  $ACR^- = \{/Employee/EmployeeRecord/profile/Pay/*, /Employee/EmployeeRecord/profile/Pay/Allowance\}$ , and  $R = \{/Employee/EmployeeRecord/profile/Pay/Allowance\}$ .  $R$  is contained in  $ACR^{-new}$ . The other negative rule prevents access to all children of  $\langle Pay \rangle$ , so the access to  $\langle Pay/Allowance \rangle$  is also prevented.  $\square$

- **[Default Rule]:** Execute the following:

$$V'(D) = V(D) \cup (R_1(D) -^D R_2(D)) \quad (5)$$

where  $R_1 = R \cap^D ACR^+$  and  $R_2 = R \cap^D ACR^- \cap^D ACR^+$  to minimize the data being brought to the view from the database.

**Example 19.** Let  $ACR^+ = \{/Employee/EmployeeRecord/profile/*\}$ ,  $ACR^- = \{/Employee/EmployeeRecord/profile/Pay/*, /Employee/EmployeeRecord/profile/Pay/Allowance\}$ , and,  $R = \{/Employee/EmployeeRecord/profile/Pay/*\}$ . In this case,  $R' = \{/Employee/EmployeeRecord/profile/Pay/*\}$  and  $R'' = \{/Employee/EmployeeRecord/profile/Pay/Allowance\}$  in Equation ???.  $\square$

A query optimizer can use the equations given above and rewrite it using different cost models to obtain equivalent equations. We show the choices available to a query optimizer for the case where a negative rule is deleted. Similar equivalent equations can be derived for the rest of the cases using set theory. We leave the rest of the cases to the reader to avoid equation clutter. In Equation ??, the sub-expression  $(R_1(D) -^D R_2(D))$  is sent to the database and the returned results are added to the view. A query optimizer can use the following equivalent sub-expressions instead of the sub-expression above if it estimates that the other sub-expressions will be computationally cheaper:

$$\begin{aligned} (R_1(D) -^D R_2(D)) &\equiv (R_1(D) -^D ACR^-(D)) \\ &\equiv (R_1(D) -^D R_3(D)) \equiv (R_1(D) -^D R_4(D)) \end{aligned}$$

where  $R_3(D) = ACR^+ \cap^D ACR^-$ ,  $R_4(D) = R \cap^D ACR^-$ .

## 5.4 Reducing Computation Cost for Incremental View Adaptation

In this section, we try to optimize some of the expressions we derived in section 5.3 for computation costs. We transform some of the expressions mentioned in section 5.3 into equivalent expressions that reduce computation costs:

1. **Addition of a positive rule  $R$ :** Reducing Deep Expects: In Equation ??,  $R'$  is essentially used to prune out parts of  $R$  that are already in the view. In cases where view adaptation time takes precedence, the expression

$$V'(D) = V(D) \cup (R(D) -^D R''(D)) \quad (6)$$

where  $R'' = R \cap^D ACR^-$  can be used instead. This eliminates the need for two operations - running  $R'$  on the database and a deep-except operation.

2. **Removal of a positive rule  $R$ :**

A. *Incremental Adaptation:* If all ACRs in  $ACR^{+new}$  can be rewritten using the existing view, then these rewritings can be evaluated on the extension of the view to adapt it.

B. *Computation Minimization Model:* While Equation ?? minimizes the data that is transmitted from the database server to the client, we can reduce the amount of computation required to adapt the view at the server size by using the Equation ???. Note that this equation may result in more data being transmitted from the database server to the client.

$$V'(D) = V(D) -^D (R(D)) -^D ACR^{+new}(D) \quad (7)$$

3. **Addition of a Negative rule  $R$ :**

To minimize the computation cost, the following equation can be executed:

$$V'(D) = V(D) - R(D) \quad (8)$$

For deleting a negative rule  $R$ , equations defined in the previous sub-section, are optimized for minimal data transfer and reduce computing cost as well.

## 6. AUXILIARY DATA

Auxiliary data is additional data that is stored with/for each security view to make view adaptation more efficient. In this work, we suggest auxiliary data that can be stored to reduce both computation as well as communication costs. Note that the auxiliary data requires overhead in terms of extra storage, however, typically the auxiliary data that we propose to keep is small and the advantages with respect to computation and communication costs significant. The different kinds of auxiliary data that can be stored are described below.

### 6.1 Rule IDs for View Trees

In the last section, we saw that, in certain cases, even when the view contained all the results that satisfy a sub-expression, e.g.,  $R(D)$ , the view adaptation algorithm had to go to the base data. For example, in Example 13, the view contained the student and employee names brought into the view by two different rules. When the rule corresponding to the student names was being deleted, the student names had to be deleted from the view. However, because the student names and the employee names could not be differentiated, the algorithm had to go to the base data to get the student names and then delete them from the view.



To avoid such visits to the database, we can maintain a table with pairs of the form:  $\langle rule-id, view-tree-root-node \rangle$ . The identifier  $rule-id$  identifies a rule in the definition of a view and is unique across all views. The entry  $view-tree-root-node$  points to the root node of each tree that is brought into the view. Note that multiple rules may bring the same tree into a view and a single rule can bring in multiple trees. In our Example 13, we would have  $rule-id$  of the rule  $/Employee/EmployeeRecord/Profile/name$  with the  $\langle name \rangle$ -subtrees brought in by that rule and the  $rule-id$  of the rule  $/Student/Profile/name$  associated with the  $\langle name \rangle$ -subtrees brought in by that rule. Now, when we want to delete the rule  $/Student/Profile/name$ , it is easy to identify the nodes associated with that rule and delete them. Thus, we avoid going to the database to adapt the view.

The overhead of maintaining the  $rule-id$  table is not significant in practice and the gains of incremental view adaptation are achieved.

Note also that in the presence of  $rule-ids$ , essentially, we can think of the views as pertaining to a set of individual views corresponding to a single positive XPath rule and an union of negative XPath rules because we can identify the data in the view corresponding to each positive rule using the  $rule-id$  table but all data that is covered by all the negative rules are not in the view.

We utilize the algorithms presented by Mandhani and Suciu [17] and Xu and Ozsoyoglu [28] to rewrite a rule (or sub-expression) using a view and use an adaptation of the ideas presented in [9] to answer a rule (or sub-expression) using multiple views.

## 6.2 Auxiliary Rule Views

In the last section, we saw sub-expressions of the form  $R_1(D) -^D R_2(D)$ . These sub-expressions were evaluated at the database. However, if there exists a rewriting of the sub-expression using the existing views, then that rewriting should be used instead of sending the query to the database to save on communication and computational costs. Note that if only  $R_2(D)$  can be obtained from the view, but  $R_1(D)$  is not, we would rather send the entire sub-expression to the database because  $R_1(D) -^D R_2(D)$  can be substantially smaller in size than  $R_1(D)$ . However, if the entire subexpression can be evaluated using the view, we would rather use the views. Typically, the existing views are much smaller than the database and thus, the evaluation of the rewritings on the extensions of the existing views will be computationally faster than going to the database. This advantage is even more pronounced when the views and the databases are not on the same machine and communication costs have to be considered.

In order to facilitate the rewriting, we propose to maintain some additional views below. Consider the case when a negative rule is removed. Then, data may need to be added to the view and typically the database has to be consulted to provide the additional data that needs to be added to the view. However, when we construct the view, if we keep an auxiliary rule view corresponding to each negative rule, we can use the corresponding auxiliary rule view to identify all nodes that are in the result of the negative rule and add them to the view (if they are not also ruled out by other negative rules), thereby saving a trip to the database.

A rule view is an intermediate view constructed using

only a single positive or negative rule and materialized. Rule views are not made available to users (otherwise it would be a security breach), but only kept internally to speed up view adaptation. The semantics of all query answering using views will remain the same whether rule views are used or not. We construct one rule view for each rule in  $ACR^+$  and  $ACR^-$ . Rule views are especially useful when the data necessary to adapt a view cannot be obtained using the view as explained above.

## 6.3 Storing Node Ids for Answering Queries using Views

Each element in an XML document can be assigned a node-ID, either by the view-adaptation program, or by an XPath/XQuery processor itself. The original node ID of each element in the original XML document can be stored along with the element in the view. This can serve several purposes.

Consider the following example, where the newly added negative rule R, is not re-writable using a view.

**Example 20.**  $ACR^+ = /Employee/EmployeeRecord/profile/Pay$ ,  $ACR^- = \{ \}$  (for simplicity), and  $R = /Employee/EmployeeRecord/profile[age < 50]/Pay$ . In this case, the view contains all the nodes that belong to the result of R but, we have to go to the database to determine which  $\langle Pay \rangle$  elements should be deleted from the view.  $\square$

If the node-ID of every  $\langle Pay \rangle$  element is stored in the view, then it has two advantages: (a) R can be executed at the base data site, and instead of communicating the resulting elements that must be removed from the view to the view site, only the node-IDs of all the  $\langle Pay \rangle$  elements in  $/Employee/EmployeeRecord/profile[age > 50]/Pay$  can be sent to the view site. The  $\langle Pay \rangle$  elements in the view with corresponding node-IDs can be removed from the view. This reduces communication cost as the size of the data in an XML element would typically be larger than the size of its node-ID. (b) Rule views containing only node-IDs can be created. This way, whenever a rule is removed, node-IDs from the rule views can be used for removing/adding from/to the view as discussed in (a) above. Storing node-IDs can also enable structural join between views [27]; a detailed discussion of structural joins between views is beyond the scope of this work.

## 7. VALIDATION

We implemented our incremental XPath view adaptation algorithms to validate our proposal experimentally. In this section, we provide details of our empirical evaluation.

### 7.1 QFilter: NFA-based query filter

In the incremental view adaptation techniques described so far, in many cases, we need to obtain the intersection between a rule to be added or deleted and  $ACR^+$  or  $ACR^-$ . To quickly compute this (deep) intersection, we use an NFA-based XML query filtering framework, QFilter [15]. We choose QFilter for our work because it is independent of the XML query processing engine and provides good performance [15] for  $XP\{/,//,*,[]\}$ . To construct a QFilter from a set of XPath expressions, an NFA for each expression is constructed and all such NFAs are merged to form a QFilter.




Figure 5: QFilter.

**Example 21.** Consider  $ACR^+ = \{\}$ ,  $ACR^- = \{/site/people/person/creditcard,/site/people/person/income,/site/regions/*/item/quantity,/site/regions/*/item/id,/site/open_auctions/open_auction/privacy\}$

The QFilter constructed from  $ACR^-$  is shown in Figure 5. To identify parts of a negative rule  $R = /site/people/person/*$  that conflict with  $ACR^-$  (i.e.  $R -^D ACR^-$ ),  $R$  is filtered through the QFilter in Figure 4. Filtering  $R$  would yield  $\{/site/people/person/creditcard \cup /site/people/person/income\}$ , i.e.,  $R' = R -^D ACR^-$  □

If predicates are involved, QFilter returns an XPath expression with conjunctions. The intersection produced is not necessarily minimized (i.e., some paths from the intersection query can be dropped and the resultant XPath query remains equivalent to the original query). However, the intersection XPath expression can be used to query the database or the view as required. Consider the following example:

**Example 22.**  $ACR^+ = /site/regions/*/item[@type="A"]/quantity$  and  $R = /site/regions/*/item[location]/quantity$ , then  $R' = /site/regions/*/item[@type="A"]/location/quantity$ . □

For more on how QFilter handles various predicates, please refer to [15].

## 7.2 Implementation Set-Up

We have used the benchmark XML schema, XMark [2], and the XMLGen document generator to generate the database of XML Documents. We used Galax [1] XQuery processor on a machine running on Linux 2.4.18, with 1 GB RAM and Intel Xeon 2.80GHz processor. The data size used for the experiments varied from 20 to 100 MB. While recording the time for our experiments, we ignored the time taken for internal data structure initializations in the Galax Package. We record the running times for each view adaptation algorithm from the start of the evaluation of the algorithm till the view has been adapted. We did not implement XPath containment and re-writing. We used the QFilter implementation and observed the performance of Equation ??-Equation ??. The view adaptation algorithms suggested in Section 5.3 are optimal algorithms that minimize the amount of incremental data added or deleted. Therefore, we provide empirical results only for an implementation of incremental view adaptation for reduced computation costs.

## 7.3 Experimental Results

We collectively refer to the complete set of access control rules defining the view i.e.  $ACR^+ \cup ACR^-$  as ACR.




Figure 6: Effect of varying number of access control rules: (a) Remove negative rule; (b) Remove positive rule; (c) Add positive rule; and (d) Add negative rule.

I. The effect of increasing the size of ACR. We studied the effect of increasing size of ACR (10 to 100) on the performance of the different view adaptation schemes. The size of the data file used was 20 MB, and the number of positive and negative rules defining the views was equal.

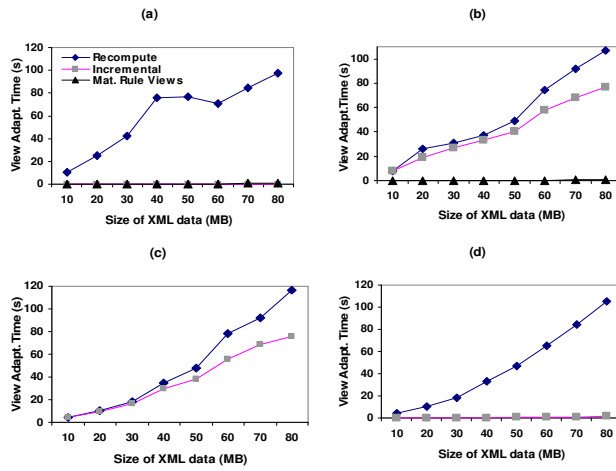
In Figures 6 and 7, we use the term "Mat. Rule Views" to represent the materialized auxiliary rule views described above. When a rule is being added, we assume that most of these rules are new rules. Therefore, the system does not have them materialized and this option is not available for the addition of positive and negative rules. However, these materialized auxiliary rule views are used when rules are removed from the views.

As shown in Figure 6, the incremental adaptation schemes performed significantly better than the re-computation scheme in all the four cases. Re-computation time increases with increase in size of  $ACR^+ / ACR^-$  because the number of XPath expressions to be evaluated during view adaptation also increases. Incremental adaptation on the other hand only needs to compute that part of the data that needs to be added/removed from the view by the new rule. As can be seen:


(a) Incremental adaptation performs best when adding a negative rule and removing a positive rule, because, for both these scenarios, in most cases, *only the view* is used for adaptation, which is typically much smaller compared to the database. (b) While removing a negative rule, incremental adaptation using auxiliary rule views performs the best. This is because, while normal incremental adaptation uses the existing view as well as the database, the auxiliary rule-view-based incremental adaptation uses the existing view and a *small auxiliary rule view* instead to adapt the view.

In Figures 6 (a-d), the sudden peak when the size of ACR increases from 60 to 70 because the new access control rules that are added bring additional data that is disproportionately large for these additional 10 new rules (which increases the time to compute all XML nodes contained by these rules from the base data).

(c) As can be seen in Figure 6, compared to the re-computation algorithm, the incremental algorithms scale much better to this change.



**Figure 7: Scalability with increasing data size: (a) Remove positive rule; (b) Remove a negative rule; (c) Add a positive rule; and (d) Add a negative rule.**



**Figure 8: Effect of increasing view size re-computation and incremental approaches: (a) Add negative rule; and (b) Add positive rule.**

**II. The effect of varying the base data size.** Figures 7(b) and 7(c) show that when the newly added positive rule (or (part of) the removed negative rule) needs to be evaluated on the base data; as the base data size increases, the time taken for evaluating an XPath expression on it also increases. In all four cases, the incremental methods perform better than re-computation. Increasing the base data size does not have any significant effect on the performance of incremental view adaptation in the cases where only the view is used for adaptation (Figures 7(a) and 7(d)). As a result: (a) In cases where only the view is used for adaptation, incremental methods perform significantly better than re-computation (sometimes by a factor of 170). (b) Where a negative rule is being removed, we see (again) that the auxiliary rule views provide a very significant performance improvement over other two methods.

**III. The effect of increasing the view size.** We observed the effect in the two representative cases, where a negative rule is added (data is removed from the view) and a positive rule is added (data is added to the view). In both cases, the incremental adaptation method scaled well to increasing view sizes whereas the time taken to re-compute the view increased with increasing size of the view (Figure 8). This is because, as the rules in ACR bring in more data into the view, more computations are needed on the base data while recomputing the view. In the incremental method, when the view size increases by small fractions, it makes very little difference to view adap-

tation time.

### 7.3.1 Summary of Observations

From the observations above, it is clear that incremental view adaptation outperforms view re-computation, in the event of all the four different kinds of updates: adding a positive rule, deleting a positive rule, adding a negative rule, and deleting a negative rule. Our incremental view adaptation approach also scales well to increasing number of access control rules, increasing sizes of base data and security views, as demonstrated by the experiments. Rule views have led to significant performance improvements for deleting positive and negative rules, wherever the rule views can be used without any rewritability or answerability issues.

## 8. CONCLUSIONS AND FUTURE WORK

Efficient view adaptation for reducing view-downtime is important for the use of access-control views in a wide range of applications. In this work, we have proposed and implemented a novel and efficient approach to incrementally adapt XPath access-control views expressed defined using  $XP\{/,//,*,[]\}$ . We also suggest techniques for more efficient view adaptation using auxiliary data, such as rule views. We show empirically that incremental view adaptation performs better than view re-computation, and that auxiliary rule views lead to significant performance improvement when positive rules are added and negative rules removed.

The solution of the XPath view adaptation problem is dependent on efficient solutions to the XPath query rewriting using multiple views problem. Therefore, solving the problem of rewriting XPath queries using XPath views especially in the presence of negative rules and plugging it into our infrastructure, will improve the efficiency and completeness of our view adaptation algorithms. The Galax XQuery processor is limited in the size of the data that it can process. A more robust XQuery processing engine will help us verify the scalability of our algorithms for even larger sizes of XML documents. Given the trends observed, the potential for gains over the recomputation-based methods is even more for larger documents. The investigation of incremental XPath view adaptation under different models of XML access control can also be undertaken.

## 9. REFERENCES

- [1] Galax xquery processor, available at <http://db.bell-labs.com/galax>.
- [2] Xmark benchmark schema, available at <http://www.xml-benchmark.org/>.
- [3] Xpath, available at <http://www.w3.org/tr/xpath20>.
- [4] A. Balmin, F. Ozcan, K. S. Beyer, R. J. Cochrane, and H. Pirahesh. A framework for using materialized xpath views in xml query processing. In *VLDB Conference*, 2004.
- [5] L. Bouganim, F. Ngoc, and P. Pucheral. Client-based access control management for xml documents. Technical report, INRIA, June 2004.
- [6] E. Damiani, S. D. C. di Vimercati, S. Paraboschi, and P. Samarati. Design and implementation of an access control processor for XML documents. *Computer Networks (Amsterdam, Netherlands: 1999)*, 33(1–6):59–75, 2000.

- [7] W. Fan, C.-Y. Chan, and M. Garofalakis. Secure xml querying with security views. In *SIGMOD '04: Proceedings of the 2004 ACM SIGMOD international conference on Management of data*, pages 587–598, New York, NY, USA, 2004. ACM Press.
- [8] I. Fundulaki and M. Marx. Specifying access control policies for xml documents with xpath. In *The ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 61–69. ACM Press, 2004.
- [9] J. Gao, D. Yang, and T. Wang. Efficient xml query rewriting over the multiple xml views. In *Proceedings of the 17th Data Engineering Workshop, (DEWS '06)*, March 2006.
- [10] G. Guerrini, M. Mesiti, and D. Rossi. Impact of xml schema evolution on valid documents. In *WIDM '05: Proceedings of the 7th annual ACM international workshop on Web information and data management*, pages 39–44, New York, NY, USA, 2005. ACM Press.
- [11] A. Gupta and I. S. Mumick. Maintenance of materialized views: Problems, techniques and applications. *IEEE Quarterly Bulletin on Data Engineering; Special Issue on Materialized Views and Data Warehousing*, 18(2):3–18, 1995.
- [12] A. Gupta, I. S. Mumick, and K. A. Ross. Adapting materialized views after redefinitions. In M. J. Carey and D. A. Schneider, editors, *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data, San Jose, California, May 22-25, 1995*, pages 211–222. ACM Press, 1995.
- [13] G. Kuper, F. Massacci, and N. Rassadko. Generalized xml security views. In *SACMAT '05: Proceedings of the tenth ACM symposium on Access control models and technologies*, pages 77–84, New York, NY, USA, 2005. ACM Press.
- [14] C.-H. Lim, S. Park, and S. H. Son. Access control of xml documents considering update operations. In *XMLSEC '03: Proceedings of the 2003 ACM workshop on XML security*, pages 49–59, New York, NY, USA, 2003. ACM Press.
- [15] B. Luo, D. Lee, W.-C. Lee, and P. Liu. Qfilter: fine-grained run-time xml access control via nfa-based query rewriting. In *The thirteenth ACM international conference on Information and knowledge management (CIKM)*, pages 543–552, New York, NY, USA, 2004. ACM Press.
- [16] B. Luo, D. Lee, W.-C. Lee, and P. Liu. Deep set operators for xquery. In *XIME-P*, 2005.
- [17] B. Mandhani and D. Suciu. Query caching and view selection for xml databases. In *VLDB '05: Proceedings of the 31st international conference on Very large data bases*, pages 469–480. VLDB Endowment, 2005.
- [18] G. Miklau and D. Suciu. Containment and equivalence for a fragment of xpath. *J. ACM*, 51(1):2–45, 2004.
- [19] M. Murata, A. Tozawa, M. Kudo, and S. Hada. Xml access control using static analysis. In *CCS '03: Proceedings of the 10th ACM conference on Computer and communications security*, pages 73–84, New York, NY, USA, 2003. ACM Press.
- [20] F. Neven and T. Schwentick. Xpath containment in the presence of disjunction, dtids, and variables. In *International Conference on Database Theory*, 2003.
- [21] R. J. Peters and M. T. Ozsu. An axiomatic model of dynamic schema evolution in objectbase systems. *ACM Trans. Database Syst.*, 22(1):75–114, 1997.
- [22] A. Sawires, J. Tatemura, O. Po, D. Agrawal, and K. S. Candan. Incremental maintenance of path-expression views. In *SIGMOD '05: Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pages 443–454, New York, NY, USA, 2005. ACM Press.
- [23] T. Schwentick. Xpath query containment. *SIGMOD Rec.*, 33(1):101–109, 2004.
- [24] M. Staudt and M. Jarke. Incremental maintenance of externally materialized views. In *The VLDB Journal*, pages 75–86, 1996.
- [25] A. Stoica and C. Farkas. Secure xml views. In *DBSec*, pages 133–146, 2002.
- [26] Y. Velegrakis, R. J. Miller, and L. Popa. Mapping adaptation under evolving schemas. In *VLDB*, pages 584–595, 2003.
- [27] Y. Wu and H. Jagadish. Structural join order selection for xml query optimization. In *ICDE Conf., Bangalore, India*, Mar. 2003.
- [28] W. Xu and Z. M. Ozsoyoglu. Rewriting xpath queries using materialized views. In *VLDB '05: Proceedings of the 31st international conference on Very large data bases*, pages 121–132. VLDB Endowment, 2005.
- [29] T. Yu, D. Srivastava, L. Lakshmanan, and H. Jagadish. Compressed accessibility map: Efficient access control for xml. In *VLDB Conference*, 2002.