Keys for XML*

Peter Buneman Susan Davidson Wenfei Fan† Carmem Hara‡
Wang-Chiew Tan

Abstract

We discuss the definition of keys for XML documents, paying particular attention to the concept of a *relative key*, which is commonly used in hierarchically structured documents.
Keywords: Keys, Relative Keys.

1 Introduction

Keys are an essential part of database design [2, 10]: they are fundamental to data models and conceptual design; they provide the means by which one tuple in a relational database may refer to another tuple; and they are important in update, for they enable us to guarantee that an update will affect precisely one tuple. More philosophically, if we think of a tuple as representing some real-world entity, the key provides an invariant connection between the tuple and entity.

If XML documents are to do double duty as databases, then we shall need keys for them. In fact, a cursory examination[1] of existing DTDs reveals a number of cases in which some element or attribute is specified – in comments – as a "unique identifier". Moreover a number of scientific databases, which are typically stored in some special-purpose hierarchical data format which is ripe for conversion to XML, have a well-organized hierarchical key structure.

Various forms of key specification for XML are to be found in the XML standard [5], XML Data [9], XML Schema [11]. Through the use of ID attributes in a DTD [5], one can uniquely identify an element within an XML document, but it is not clear that ID attributes are intended as keys rather than internal "pointers" For example, ID attributes are not scoped. In contrast to keys, they are unique within the entire document rather than among a designated set elements. As a result, one cannot, for example, allow that a student (element) and a person (element) use the same SSN as an ID. Moreover using ID attributes as keys means that we are limiting ourselves to unary keys and, of

*University of Pennsylvania. Email: {peter,susan,wctan}@saul.cis.upenn.edu

†Temple University. Email: fan@joda.cis.temple.edu. Supported by RIF fund, Temple University.

‡Universidade Federal do Parana, Brazil. Email: carmem@inf.ufpr.br

[1]We used the "DTD Inquisitor" of Byron Choi and Arnaud Sahuguet [4, 7].

course, to using attributes rather than elements. Finally, one can specify at most one ID attribute for an element type, while in practice one may want more than one key. XML Data introduces a notion of keys explicitly. However, its keys can only be specified in types and moreover, can only be defined for element types rather than for certain collections of elements.

XML Schema has a more elaborate proposal, which is the starting point of this paper. The proposal extends the key specification of XML Data by allowing one to specify keys in terms of XPath [8] expressions. There are a number of technical problems in connection with XPath. XPath is a relatively complex language in which one can not only move down the document tree, but also sideways or upwards, not to mention that predicates can be embedded as well. The main problem with XPath is that questions about equivalence or inclusion of XPath expressions are, as far as the authors are aware, unresolved; and these issues are important if we want to reason about keys as we do in relational databases. Yet until we know how to determine the equivalence of XPath expressions, there is no general method of saying whether two such specifications are equivalent. Another technical issue is value equality. XML Schema restricts equality to text, but the authors have encountered cases in which keys are not so restricted. A more detailed discussion can be found in section 6.1.

However, the main reason for writing this note is that none of the existing key proposals address the issue of hierarchical keys, which appear to be ubiquitous in hierarchically structured databases, especially in scientific data formats. A top-level key may be used to identify components of a document, and within each component a secondary key is used to identify sub-components, and so on. Moreover, the authors believe that the use of keys for citing parts of a document is sufficiently important that it is appropriate to consider key specification independently of other proposals for constraining the structure of XML documents.

How then, are we to describe keys for XML or, more generally, for semistructured data? From the start, how we identify components of XML documents is very different from the way we identify components of relational databases. Consider the two structures:

```
⟨db⟩
   ⟨student⟩ ⟨name⟩ Smith ⟨/name⟩ ⟨course⟩ Math2 ⟨/course⟩ ⟨grade⟩ B ⟨/grade⟩ ⟨/student⟩
   ⟨student⟩ ⟨name⟩ Jones ⟨/name⟩ ⟨course⟩ Math2 ⟨/course⟩ ⟨grade⟩ A+ ⟨/grade⟩ ⟨/student⟩
   ⟨student⟩ ⟨name⟩ Brown ⟨/name⟩ ⟨course⟩ Phil5 ⟨/course⟩ ⟨grade⟩ A- ⟨/grade⟩ ⟨/student⟩
⟨/db⟩
```

name	course	grade
Smith	Math2	B
Jones	Math2	A+
Brown	Phil5	A-

To identify a tuple in the relation we need to know, say, that `name` and `course` constitute a key. In

the absence of a key the only way we can be sure of uniquely identifying a tuple is to give the entire tuple. For relational databases, the way we specify a key constraint is to say that if two tuples agree on their key attributes they agree everywhere. By contrast, XML documents are, first of all, documents and we can therefore use the position in the document (say a byte offset) to identify some part of it. However, when faced with the problem of updates, it is best to use value-based identifiers rather than location-based identifiers, e.g. we would like to be able to say that if two elements agree on the `name` and `course` subelements then they are the same element. Put in the contrapositive: two distinct `student` elements must differ on a `name` or `course` subelement. This raises two issues that precede any discussion of the structure of keys: that of node identification and that of equality. The latter is a thorny topic, but needs some attention.

Organization. The rest of the paper is organized as follows. Section 2 introduces the notion of node addresses and value equality. Node addresses are used in node equality testing, ie, testing whether two nodes are the same node and value equality is used for testing whether two nodes have the same value. Section 3 introduces our path expression language which is used in the definition of keys discussed in section 4. The concept of relative or hierarchical keys together with its alternative notation is discussed in section 5. In section 6, we examine the XML-Schema proposal in some detail, discuss an an alternative form of keys and various issues concerning keys.

2 Node addresses and equality

The *Document Object Model* (DOM) [3] provides some insight into a semantics for XML documents. According to the DOM, a document is a hierarchical structure of nodes. Nodes are of several types, but there are three types that are important to this discussion: element nodes, attribute nodes, and text nodes. As illustrated in Figure 1 text nodes (T) have no name but carry text, attribute nodes (A) both have a name and carry text, and element nodes (E) have a name. Element nodes may have children; attribute and text nodes are terminal. In addition the DOM specifies how to reach the children of an element node. Text and element children are held in what is essentially an array, the index in the array being determined by the order of the subelements in the document. Attribute children are held in a dictionary. The name of the attribute, which must be unique within an element, is used as the index. These indexes, an integer for an element or text child, or the name prefixed by an "@" for attributes, are shown as edge labels in Figure 1. The important point here is that the edge labels uniquely identify children.

A consequence of this model is that a path of edge labels from the root uniquely identifies a node. We shall call such paths *node addresses* and write them $\langle l_1 \# \ldots \# l_n \rangle$, for example $\langle$`1#2#1`$\rangle$ and $\langle$`1#3#@num`$\rangle$. Node addresses will be our basic means of identifying nodes. Note that an attribute name can only occur at the end of a node address. We can also talk about the address of a subnode

```
⟨db⟩
   ⟨composer⟩
      ⟨name⟩ J.S. Bach ⟨/name⟩
      ⟨born⟩ 1685 ⟨/born⟩
      ⟨work  num="BWV82"⟩
         ⟨title⟩ Ich habe genug ⟨/title⟩
      ⟨/work⟩
      ⟨work  num="BWV552"⟩
      ⟨/work⟩
   ⟨/composer⟩
   ⟨composer  period="baroque"⟩
      ⟨name⟩ G.F. Handel ⟨/name⟩
      ⟨work  num="HWV19"⟩
         ⟨title⟩ Art Thou Troubled ⟨/title⟩
      ⟨/work⟩
   ⟨/composer⟩
⟨/db⟩
```

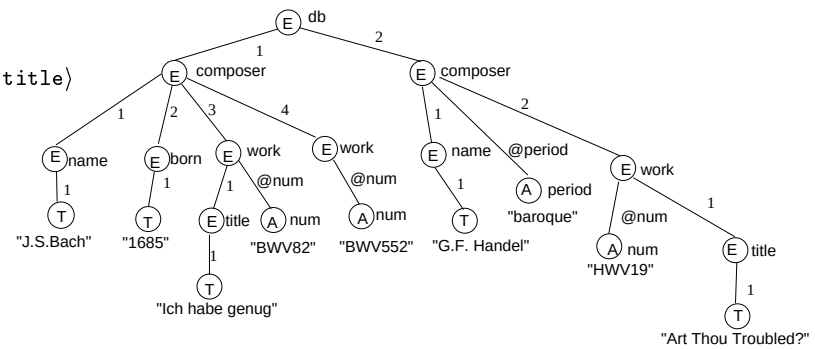

Figure 1: Some XML and its representation as a tree

relative to a node. For example any subnode of a node with address $\langle a \rangle$ will have a node address of the form $\langle a\#b \rangle$ where $\langle b \rangle$ is the address of the subnode relative to $\langle a \rangle$. By a *subnode* of a node x we mean any node in the subtree rooted at x, not necessarily a child node of x.

Value equality. Equality is essential to the definition of keys, and in order to define keys we need first to define equality of the "values" associated with nodes. XML-Schema restricts equality to text nodes, but the authors have encountered cases in which keys are not so restricted. A more general way of describing equality is to use tree equality. The value of a node is specified by giving (1) a set S of relative addresses of its subnodes, (2) a partial function from S to names and (3) a partial function from S to strings. Two nodes are *value-equal* if they agree on (1), (2) and (3). With respect to the textual representation of an XML element, this definition states that the order of attributes is unimportant in defining equality. Observe that the order of subelements is specified and preserved by their indexes (integers).

Notation. We shall use $=_v$ for value equality.

It should be pointed out that neither equality of text nodes nor tree equality is entirely satisfactory in the presence of types. XML-Schema does a thorough job of defining base types, and one might want to use this to define a coarser form of equality. For example, ⟨`id  type="int"`⟩ `0` ⟨`/id`⟩ and ⟨`id  type="int"`⟩ `-0` ⟨`/id`⟩ should probably be treated as value-equal. Also, there are types such as real numbers for which equality is problematic. A complete specification of keys would have to take account of these issues.

3 Path Expressions

A path expression is an expression involving node names (tags and attribute names) that describes a set of paths in the document tree.

The choice of what language we use to define path expressions is important to the expressive power of keys, and there are a number of choices. In XML-Schema, XPath [8] expressions are used, while in semistructured data regular expressions [1] have been used. Neither subsumes the other. In the following analysis we shall assume two properties of path expressions:

- There should be a concatenation operation: $P.Q$ is the result of following first the path P and then the path Q.
- A path should move down the tree. That is if we start at a node n_1 and, by following a path described by P, we reach a node n_2 then n_2 is a subnode of n_1 (the address n_1 is a prefix of the address n_2.)

The second property is not enjoyed by XPath. We shall discuss the choice of a language of path expressions later, but in the meantime adopt for illustrative purposes a simple language that is certainly a subset of both XPath and regular expressions. Our language, which uses a fragment of XPath syntax, has the following constructs:

- The empty path, "ϵ".
- A node name (a tag or attribute name).
- A wild card "_", matching any single node name.
- An arbitrary path "_*".
- The concatenation of paths $P.Q$, where P and Q are paths defined by these rules.

We note that there is an easy translation from our path language to XPath syntax. Any path meant to start from the root is prefixed with "/". In XPath, "/" itself denotes the root node. "." is used as the empty path in place of "ϵ", "*" in place of "_" and "//" in place of "_*". Also, "/" is used as the path concatenator in place of ".". In XPath, "/" is used as a separator between location steps. Therefore, we have to disallow certain concatenations. If for example we concatenate a/b with $/c/d$ we get $a/b//c/d$ with an entirely different meaning.

We shall use the notation $n[\![P]\!]$ to denote the set of nodes (node addresses) reached by starting at node n and following a path that conforms to (is in the language of) P. We shall sometimes use $[\![P]\!]$ as an abbreviation for $root[\![P]\!]$. The syntax is borrowed from Wadler's [12] description of

semantics for patterns in XSL. Examples (from Figure 1):

$$\begin{aligned}
\langle 2\#2\rangle[\![\mathsf{title}]\!] &= \{\langle 2\#2\#1\rangle\} \\
[\![\mathsf{composer._}]\!] &= \{\langle 1\#1\rangle, \langle 1\#2\rangle, \langle 1\#3\rangle, \langle 1\#4\rangle, \langle 2\#1\rangle, \langle 2\#2\rangle\} \\
\langle 2\#2\rangle[\![_*]\!] &= \{\langle 2\#2\rangle, \langle 2\#2\#1\rangle, \langle 2\#2\#1\#1\rangle, \langle 2\#2\#@\mathsf{num}\rangle\} \\
[\![\mathsf{composer.work}]\!] &= \{\langle 1\#3\rangle, \langle 1\#4\rangle, \langle 2\#2\rangle\} \\
[\![_*.\mathsf{num}]\!] &= \{\langle 1\#3\#@\mathsf{num}\rangle, \langle 1\#4\#@\mathsf{num}\rangle, \langle 2\#2\#@\mathsf{num}\rangle\}
\end{aligned}$$

In some cases, it will be useful to restrict the path expression language so that paths are merely sequences of labels and do not contain _ or _*. Such paths are called *simple paths*. For example, composer.work is a simple path.

4 Definition of Keys

In defining a key we specify two things: a set on which we are defining the key (in relational databases this is a relation – the set of tuples identified by a relation name) and the "attributes" (relational terminology for a set of column names) which together uniquely identify elements in the set. This is the motivation for our central definition of a *key specification*, which is a pair $(Q, \{P_1, \ldots, P_n\})$ where Q is a path expression and $\{P_1, \ldots, P_n\}$ is set of simple path expressions. The idea is that the path expression Q identifies a set of nodes, which we refer to as the *target set*, on which the key constraint is to hold. Let us refer to Q as the *target path*, and the set $\{P_1, \ldots, P_n\}$ as the *key paths*. These correspond to the absolute and relative location paths respectively in XPath terminology. Observe that for any node $n \in [\![Q]\!]$ there is a set of nodes $n[\![P_i]\!]$ found by following P_i from n. There is no restriction on the size of $n[\![P_i]\!]$; in particular it may be empty. The key paths constrain the target set as follows: Take any two nodes $(n_1, n_2) \in [\![Q]\!]$ and consider the pair of sets of nodes found by following the key path P_i from n_1 and n_2, $(n_1[\![P_i]\!], n_2[\![P_i]\!])$. If there is a non-empty intersection with respect to value equality for all such pairs of sets of nodes then the nodes n_1 and n_2 are the same node.

For example, consider the following key definition:

$$(\mathsf{person.employees}, \{\mathsf{name.firstname}, \mathsf{name.lastname}\})$$

The target path person.employees identifies a set of nodes in the document (there may be several employees and person nodes). This is the target set. Each of these nodes will define a subtree with an employees label at the root. Within a subtree we will find zero or more key paths name.firstname and zero or more key paths name.lastname. Two nodes n_1, n_2 in the target set are distinct if either

they do not agree on any of the nodes reachable via name.firstname or they do not agree on any of the nodes reachable via name.lastname.

As another example, observe that the document in Figure 1 satisfies the key (composer, {name}): There are two nodes at the end of the target path composer. For each node, there is one element in the set of nodes found by following the key path name, "J.S.Bach" and "G.F.Handel". These elements are not value-equal. Less intuitively, the document also satisfies the key (composer, {born}) since the subelement <born> only appears in the first composer and is absent from the second composer.

We are now ready to give the formal definition of a key. For reasons which will emerge shortly, it is useful to define a key with respect to a given node in the document rather than assuming that the target path starts at the root.

Definition. A node n *satisfies* a key specification $(Q, \{P_1, \ldots, P_k\})$ iff for any n_1, n_2 in $n[\![Q]\!]$, if for all i, $1 \leq i \leq k$, there exist $z_1 \in n_1[\![P_i]\!]$ and $z_2 \in n_2[\![P_i]\!]$ such that $z_1 =_v z_2$, then $n_1 = n_2$. That is,

$$\forall\, n_1\, n_2 \in n[\![Q]\!]\ ((\bigwedge_{1 \leq i \leq k} \exists\, z_1 \in n_1[\![P_i]\!]\ \exists z_2 \in n_2[\![P_i]\!]\ (z_1 =_v z_2)) \rightarrow n_1 = n_2).$$

Note that both forms of equality are used in the definition of a key. The first deals with value-equality ($=_v$) while the second is node equality ($=$). Two nodes are node equal if they have the same node address.

When we talk about document satisfying a key specification we mean that the root of the document satisfies the key specification. The key has no impact on those nodes at which some key path is *missing*, i.e. nodes n such that $n[\![P_i]\!]$ is empty for some P_i. Observe that for any n_1, n_2 in $[\![Q]\!]$, if P_i is missing at either n_1 or n_2 then $n_1[\![P_i]\!]$ and $n_2[\![P_i]\!]$ are by definition disjoint. This is similar to *unique constraints* introduced in XML-Schema. In contrast to unique constraints, however, our notion of key specification is capable of comparing nodes at which a key path may lead to multiple nodes. As an example, consider a key (A, {B}) expressed with respect to the root of the following document:

```
⟨db⟩
     ⟨A⟩  ⟨B⟩ 1 ⟨/B⟩  ⟨/A⟩
     ⟨A⟩  ⟨B⟩ 1 ⟨/B⟩  ⟨B⟩ 2 ⟨/B⟩  ⟨/A⟩
⟨/db⟩
```

This key asserts that an A element is uniquely identified by the values of its B subelements. The document does not satisfy the key because the B subelement in the first A element and the first B subelement of the second A element have the same value. Thus with our notion of keys we cannot

distinguish these two A elements. In contrast, key and unique constraints of XML-Schema cannot compare these A elements because the second A element has more than one B subelement.

Here are some further examples of keys, expressed with respect to the root of a document.

$(_*.\mathsf{person}, \{\mathsf{id}\})$	Any person element, if it has id subelements, is uniquely identified by the values of the id's. In other words, any two person elements are disjoint on their id fields up to value-equality.
$(\mathsf{person}, \{\epsilon\})$	Any two person nodes immediately under the root have different values (. is the empty path).
$(\mathsf{employees}, \{\})$	An empty key. This means that the path employees, if it exists, is unique at the root. I.e., there is at most one employees node immediately under the root.
$(_*, \{\mathsf{id}\})$	Any element that has id subelements is uniquely identified by the values of the id's. That is, any two nodes are disjoint on their id fields up to value-equality. Note that an id element does not have to have an id itself. This key constraint captures the semantics of an ID attribute in the XML standard.

Along the same lines of keys in relational databases, this notion of keys asserts that the values associated with key paths uniquely identify a node in the target set. However since one cannot require XML documents to be in some kind of first normal form, there are several salient differences between this notion of key and that of relational databases [2, 10]. First, the paths that define keys need not exist and do not have to be unique.[2] In contrast, in relational databases key values cannot be null (the key must exist) and first normal form requires attribute values to be individuals, not sets. Second, our keys specify a set of *addresses* within a document, unlike the relational case in which keys specify a value.

There are, of course, other ways of defining keys, both more and less restrictive than what we have described. Some justification of the choices is in order.

- We have used a *set* of key paths to define a key. In order to talk about a set (as opposed to a tuple or list) of path expressions we need to be able to talk about equality of path expressions. The equivalence of two path expressions in our language of path expressions is decidable, as it is for the more general class of regular expressions.

- Given that we have defined equality on trees, do we need to have more than one key path in a key specification? We could always design our documents so that all the key "attributes" are represented as subnodes of some node. The problem here is that we would have to constrain the node to contain only these subnodes for tree equality to have the desired effect. This seems

[2] This might be taken as allowing *null-valued* keys, but whether we should equate missing key paths with null values is arguable and depends on the semantics of the languages we use to query XML documents.

to be too restrictive and constitutes unnecessary interference between key specifications and data models.

- The definition of key satisfaction differs significantly from the relational case by allowing a (possibly empty) set of nodes at the end of each key path. We shall examine a more restrictive definition in which key satisfaction requires each of the key paths to exist and to be unique from any node in $n[\![Q]\!]$ in Section 6.

- The language of path expressions may be regarded both as too weak and too powerful. Consider the key $(Q, \{P_1, \ldots, P_k\})$: For now, we have allowed Q to be an arbitrary path expression but have restricted the P_i to be simple paths. Would one ever want an arbitrary path $(_*)$ in one of the P_i? Also, it is not hard to come up with examples in which one would like something more powerful to express Q, e.g., $(\mathsf{person}.(\mathsf{mother}|\mathsf{father})*, \{\mathsf{id}\})$. This means a person element followed by zero or more father or mother elements. Our emphasis is that the language of path expressions is provisional, and that allowing arbitrary path expression for the P_i merely complicates the definition of key but does not change much in the way of the theory.

As in relational databases we can infer some keys from the presence of others. In particular, the notion of a "superkey" still holds:

Fact. If $(Q,\ S)$ and $S \subseteq S'$, then $(Q,\ S')$.

In contrast to relational databases, a number of other inference rules also hold. Two examples are given below; the axiomatization of such inference rules and the complexity of reasoning about keys are the topic of another paper [6].

Fact. If $(Q.Q', \{P\})$ then $(Q, \{Q'.P\})$.

Fact. If $(Q, S \cup \{P_i, P_j\})$ and P_i is contained in P_j (i.e., the path language defined by P_i is included in the one defined by P_j), then $(Q, S \cup \{P_i\})$.

This last fact illustrates why we care about containment of path expressions.

Another natural question to ask is whether a key constraint is finitely satisfiable. In relational databases, all keys are finitely satisfiable: given any schema S and any finite set Σ of keys, one can always construct a finite database instance of S that satisfies Σ. The same holds for XML documents under our definition of key since keys can be missing (recall the $(_*, \{\mathsf{id}\})$ example):

Fact. For any finite set Σ of keys, there exists an (finite) XML document satisfying Σ.

This property holds in the absence of DTDs. It should be mentioned that there is interaction between DTDs and keys, but this is the topic of another paper.

5 Relative Keys

There are many situations in which a key provides only a relative specification. For example, a verse number is unique only within a chapter. In relational database design, the key of a weak entity is made up of the key of the "parent" entity and some additional identification [10], e.g. course Math120, section B. To describe this we need the notion of a *relative* key, which consists of a pair (Q, K) where Q is a path expression and K is a key.

Definition. A document satisfies a relative key specification (Q, (Q',S)) iff for all nodes n in $[\![Q]\!]$, n satisfies the key (Q', S).

In other words (Q, K) is a relative key if K is a key for every "sub-document" rooted at a node in $[\![Q]\!]$. Examples:

(bible.book.chapter, (verse, {number}))	A verse number uniquely identifies a verse within a chapter.
(bible.book, (chapter, {number}))	Chapter numbers uniquely identify a chapter within a book.
(bible, (book, {name}))	If there is only one bible node immediately under the root, this is the same as specifying an absolute key (bible.book, {name}).

Observe that in a relative key $(Q, (Q', S))$, Q is from the root whereas Q' starts at a node in $[\![Q]\!]$. It is for this reason that we defined key satisfaction at arbitrary nodes.

Transitivity of relative keys. The purpose of keys is to specify uniquely certain components of a document. Obviously, a relative key such as (bible.book.chapter, (verse, {number})) alone does not uniquely identify a particular verse in the bible, however we believe that if we give a book name, a chapter number, and a verse number, we have specified a verse. It is this intuition that we need to formalize.

First observe that the relative key $(\epsilon, (Q', S))$ is equivalent to the key (Q', S). Now consider two relative keys. We say that $(Q_1, (Q'_1, S_1))$ *immediately precedes* $(Q_2, (Q'_2, S_2))$ if $Q_2 = Q_1.Q'_1$. Any relative key immediately precedes itself. Define the *precedes* relation as the transitive closure of the immediately precedes relation.

Definition. A set Σ of relative keys is *transitive* if for any relative key $(Q_1, (Q'_1, S_1)) \in \Sigma$ there is a key $(\epsilon, (Q'_2, S_2)) \in \Sigma$ which precedes $(Q_1, (Q'_1, S_1))$.

As an example, this set of keys is transitive:

$$(\epsilon, (\mathsf{bible.book}, \{\mathsf{name}\}))$$
$$(\mathsf{bible.book}, (\mathsf{chapter}, \{\mathsf{number}\}))$$

This set is not:

$$(\epsilon, (\mathsf{bible.book}, \{\mathsf{name}\}))$$
$$(\mathsf{bible.book.chapter}, (\mathsf{verse}, \{\mathsf{number}\}))$$

Observe that keys are special cases of relative keys since they are expressible in the form of $(\epsilon, (Q, \{P_1, \ldots, P_k\}))$. Any transitive set of relative keys must contain some key.

Updatable relative keys. Consider the following (transitive) key specification:

$$(\epsilon, (\mathsf{university}, \{\mathsf{name}\}))$$
$$(\mathsf{university}, (\mathsf{dept.employee}, \{\mathsf{emp\text{-}id}\}))$$

To *identify* an `employee` node in this database, we need only to specify a university name and an `emp-id` within that university. To *add* a new employee to the database, we clearly need to specify a department for the employee. However, although this key specification is transitive, there is no way to identify a department and hence no way to add an employee. This motivates our final definition of *updatability* as shown below: With updatability, one can always insert an element in the "keyed" part of the document unambiguously by specifying (using keys) where to insert.

Definition. A set Σ of relative keys is *updatable* if it is transitive and whenever $(Q_1, (Q_2.n, S_1)) \in \Sigma$ there is a relative key $(Q_1, (Q_2, S_2)) \in \Sigma$ where $|Q_2| > 0$. Here n is a node name.

Informally, this definition gives us the property that every element with a prefix along the path $Q_1.Q_2$ can be identified through some keys. Therefore, it is easy to see that the addition of the following key will make the previous example updatable. In particular, to insert an employee, we now can specify which department (in addition to the university).

$$(\mathsf{university}, (\mathsf{dept}, \{\mathsf{dept\text{-}name}\}))$$

Even though we can now add new employees, there is still something anomalous: Although employees are nested under departments, nothing about the department is necessary to identify them. This is reminiscent of the anomalies that occur in non-second normal form relational databases. There is something wrong with the design of this document in that employees should not be children of department nodes, but only of university nodes. The linkage between employees and departments should be expressed through a foreign key. Formalizing the concept of a well-designed document with respect to its key specification is beyond the scope of this paper.

5.1 A notation for relative keys

If a system of relative keys is transitive, it forms a hierarchical structure. We can therefore create a compressed syntax for such systems.

The basic syntactic form is

$$Q_1\{P_1^1, \ldots, P_{k_1}^1\}.Q_2\{P_1^2, \ldots, P_{k_2}^2\}. \ \ldots \ .Q_n\{P_1^n, \ldots, P_{k_n}^n\}$$

This describes a system of relative keys. For each i, $1 \leq i \leq n$, it defines the relative key $(Q_1. \ \ldots \ .Q_{i-1}, (Q_i, \{P_1^i, \ldots, P_{k_i}^i\}))$. It should be noted that the first of these is of the form $(\epsilon, (Q_1, \{P_1^1, \ldots, P_{k_1}^1\}))$ and is a key.

For example bible{}.book{name}.chapter{number}.verse{number} specifies the updatable system of keys:

$(\epsilon, (\text{bible}, \{\}))$
$(\text{bible}, (\text{book}, \{\text{name}\}))$
$(\text{bible.book}, (\text{chapter}, \{\text{number}\}))$
$(\text{bible.book.chapter}, (\text{verse}, \{\text{number}\}))$

So far the key hierarchies we have specified are linear. Consider the following two specifications: company{name}.employee{id} and company{name}.department{name}. It is helpful to fold these into a single specification:

company{name}[.employee{id}, .department{name}]

This is simply a syntactic shorthand: $R[R_1, \ldots, R_n]$ for $RR_1, \ldots, RR_n$.

As a further example, consider

university{name}.school[{name}, .department[{name}, .student{id}]]

This is another example of an updatable set of relative keys. Note that students are uniquely specified by an id within a university (not just within a department or school)

Specifications such as these are reasonably compact and understandable. Their importance is not only to ensure the internal consistency of a document, but also to tell others how to cite a component of our document. This is especially important if the document is subject to change. Even though we have constructed a minimal system for describing hierarchical key structures, it turns out that this takes us some way towards describing a data model. Contrast relational database specification student(snum, name, major) and enroll(snum,cnum,grade) with a key specification

[student{snum}[.name{}, .major{}], enroll{snum,cnum}.grade{}]

They describe closely related structures. The specification [name{}, major{}] ensures that under a student node there is at most one name and at most one major node. However the key specification allows other unspecified nodes to occur under a student node and, of course, it does not require any kind of first normal form. Nevertheless, we can specify that our documents have a structured "core" somewhat akin to the complex object or nested relational structures that have been studied in databases [2]. Not surprisingly there is close interaction between key constraints and data models which requires much further study.

6 Discussion

Our main reason for writing this document was to clarify the notion of a relative key and to understand the hierarchical key structure that appears to occur naturally in a variety of data formats. What we have described here is a proposal for a key definition, and there are a number of variations on this definition which should be considered. This section contains a brief review of those alternatives, starting with the proposals in XML-Schema.

6.1 XML-Schema

XML-Schema includes a syntax for specifying keys which is related to our definition, but there are some substantive differences, even if we ignore the issue of relative keys. Possibly the most important of these is that the language for path expressions is XPath. XPath is a relatively complex language in which one can not only move down the document tree, but also sideways or upwards, not to mention that predicates can be embedded as well. The main problem with XPath is that questions about equivalence or inclusion of XPath expressions are, as far as the authors are aware, unresolved; and these issues are important if we want to reason about keys as we do – for quite practical purposes – in relational databases. Here is a brief summary of the other salient differences between our definitions and the XML-Schema proposal.

Equality. We have used a more general form of equality than that in XML-Schema. However, as pointed out in Section 2 a full treatment of equality might involve types or even some form of user-defined equality.

Definition of the target set. In XML-schema the path expression that defines the target set is taken to start at arbitrary nodes. Thus if Q defined the target set in XML-schema, this is equivalent to $_*.Q$ in our notation, starting from the root. It is, of course, possible to "root" a path expression XML-Schema.

Definition of key paths. XML-Schema talks about a *list* (not a set) of key paths. While this avoids issues of equivalence of XPath expressions, one can construct keys that are, presumably, equivalent, but have different or anomalous presentations. For example (temporarily using [...] for lists): (person,[firstname, lastname]), (person,[lastname, firstname]), (person,[lastname, lastname, firstname]) impose the same constraint. Since the issue of equivalence of XPath expressions is unresolved, there is no general method of checking whether two such specifications are equivalent.

Relative keys. While there is no direct notion of a relative key in XML-Schema, in certain circumstances one can achieve a related effect. Consider for example:

book{name}.chapter{name}.verse{number}.

In XML-Schema one can specify a key for verse (this is not XML-Schema syntax) as

(book.chapter.verse, [number, *up*.name, *up*.*up*.name]).

Here "*up*" is the XPath instruction to move up one node. Thus part of the key is outside of the value of a verse node. One of the inferences one could make for such a specification is that (book.chapter, [name, *up*.name]) is a key *provided* the nodes in the target set all contain at least one verse child node. Again, it is not clear how to reason generally about such specifications.

6.2 Some stronger definitions of keys

The definition of keys we have adopted in this paper is quite weak, which we believe is in keeping with the semi-structured nature of XML. However, intuitively it does not mirror the requirements imposed by a key in relational databases, i.e. the *uniqueness* of a key and *equality* of key values. We now explore a definition which captures both these requirements.

Strong keys. In a strong key definition, we require that the keys paths exist and are *unique*, i.e. $n[\![P_i]\!]$ contains exactly one node for $1 \leq i \leq n$. The key paths constrain the target set as follows: Take any two nodes $(n_1, n_2) \in [\![Q]\!]$ and consider the pairs of nodes found by following a key path P_i from n_1 and n_2. If all such pairs of nodes are value-equal, then the nodes n_1 and n_2 are the same node.

As an example of what it means for a path expression to be unique, consider Figure 1: name is unique at $\langle 1 \rangle$, but work and num are not unique at this node.

The definition of satisfaction for strong keys now becomes the following.

Definition. A node n *satisfies* a key specification $(Q, \{P_1, \ldots, P_k\})$ if

- For all n' in $n[\![Q]\!]$ and for all $P_i (1 \leq i \leq k)$, P_i is unique at n'.
- For any n_1, n_2 in $n[\![Q]\!]$, if $n_1[\![P_i]\!] =_v n_2[\![P_i]\!] (1 \leq i \leq k)$ then $n_1 = n_2$.

Given this strong notion of keys, let us re-examine some examples given before.

$(_*.\mathsf{person}, \{\mathsf{id}\})$	Any two person elements, no matter where they occur, have unique id subelements and differ on those elements.
$(\mathsf{person}, \{\epsilon\})$	The interpretation of this key remains unchanged under a strong key semantics.
$(\mathsf{employees}, \{\})$	Again, the semantics of this key is the same with respect to the strong and weak key specifications.
$(_*, \{\mathsf{k}\})$	This requires that every element have a key k, including any element whose name is k .

The last example points out that under a strong key semantics, finite satisfiability does not hold for all keys: $(_*, \{\mathsf{k}\})$ imposes an infinite chain of k nodes and therefore, there is no finite document satisfying it. The problem arises because we require that key paths must exist. It should be mentioned that the corresponding key in XML-Schema, $(//*, [id])$, is not meaningful either, because an id node cannot have a base type if it is to have an id subelement itself.

The concept of relative keys can be naturally adapted for strong keys as well. We say a document satisfies a strong relative key specification $(Q, (Q', S))$ iff for all nodes n in $[\![Q]\!]$, n satisfies the strong key (Q', S).

The strong notion and weak notion of keys impose different restrictions on key paths. At one end of the spectrum, all key paths must exist and be unique (strong keys). At the other end, no structural constraints are imposed on key paths (weak keys). There are also possibilities in between; for example, adopting a slightly stronger notion of weak keys which substitutes *equality* for *value intersection* of the node sets reachable by a simple key path.

6.3 Choice of a path expression language

We have used a language for path expressions that contains just enough to illustrate most of the issues that occur in connection with keys for XML. In order to reason about keys, it is essential that equivalence and inclusion of path expressions are decidable, but this is the case for the more expressive language of regular expressions, and we could equally well have used this language; none of the results would be affected. However the examples we found that used the added expressive power were somewhat contrived, and it is not clear whether this larger language is of practical use.

An interesting issue is whether, in defining a key $(Q, \{P_1, \ldots, P_n\})$, the language used to describe the target path Q needs to be the same as the language used to define the key paths $P_1, \ldots, P_n$. One could choose a *simpler* language for key paths that is a sublanguage of the language for target paths. In fact, we only require that the composition $Q.P_i$ of a target path and a key path should be in the language of target paths.

To simplify the discussion, so far we have required key paths to be simple paths. However, we could see no other benefit to simplifying the language of key paths. Below we extend the current proposal by allowing key paths to include _ and _*, i.e., to be expressed in the same language that defines target paths. To do so, we first define a notion of *value intersection*. Observe that the regular language defined by a path expression is a set of simple paths. Let us use ρ to range over simple paths. Given a path expression P, we use $\rho \in P$ to denote the simple path ρ in the language defined by P.

Value intersection. Let n_1 and n_2 be two nodes in an XML tree T and P be a path expression in the language defined in Section 3. The *value intersection* of $n_1[\![P]\!]$ and $n_2[\![P]\!]$, denoted by $n_1[\![P]\!] \cap_v n_2[\![P]\!]$, is defined as follows:

$$n_1[\![P]\!] \cap_v n_2[\![P]\!] = \{(z, z') \mid \exists\, \rho \in P, z \in n_1[\![\rho]\!], z' \in n_2[\![\rho]\!], z =_v z'\}.$$

Intuitively, $n_1[\![P]\!] \cap_v n_2[\![P]\!]$ consists of pairs of nodes that are value equal and are reachable by following the same simple path in the language defined by P starting from n_1 and n_2, respectively.

Using this notation, we extend our key specification as follows.

Key specification. A key is a pair $(Q, \{P_1, \ldots, P_n\})$, where Q and P_i's are path expressions in the language defined in Section 3. A node n *satisfies* the key iff for any n_1, n_2 in $n[\![Q]\!]$, if for all i, $1 \leq i \leq k$, the value intersection of $n_1[\![P_i]\!]$ and $n_2[\![P_i]\!]$ is not empty, then $n_1 = n_2$. That is,

$$\forall\, n_1\, n_2 \in n[\![Q]\!]\ ((\bigwedge_{1 \leq i \leq k} (n_1[\![P_i]\!] \cap_v n_2[\![P_i]\!] \neq \emptyset) \rightarrow n_1 = n_2).$$

6.4 Node names as key values

The choice of an appropriate definition for keys for XML will ultimately be determined by practice. The aim of setting out a key specification is to cover the practical cases without using definitions that are too complex to allow any kind of reasoning about keys. Have the proposals in this paper covered the practical cases?

There is one issue that may arise in "unconstrained" XML. Consider the database

```
⟨db⟩
   ⟨parts⟩
      ⟨widget⟩
         ⟨id⟩ 123 ⟨/id⟩⟨weight⟩ 1.5 ⟨/weight⟩
      ⟨/widget⟩
      ⟨widget⟩
         ⟨id⟩ 234 ⟨/id⟩⟨weight⟩ 2.5 ⟨/weight⟩
      ⟨/widget⟩
      ⟨gadget⟩
         ⟨id⟩ 123 ⟨/id⟩⟨weight⟩ 3.2 ⟨/weight⟩
      ⟨/gadget⟩
   ⟨/parts⟩
⟨/db⟩
```

The type of a part – `widget` or `gadget` – is expressed in the tag. In alternative XML representations it might be expressed as an attribute or subelement of a `part` element. The key for a part is to be taken as its type together with its `id`. With our current machinery, the key constraint can be expressed as parts{}[.widget{id}, .gadget{id}]. However, if we introduce a new part type, a `thingy`, the key specification will have to be changed to include a key path involving `thingy`. No change would be needed in the alternative representations. The problem arises because we are interchanging structure (the names) with data (their values); but the ability to do this is supposed to be one of the strong points of semistructured data and XML.

Our definition of a key (weak or strong) can be extended to express this by adding a "virtual" subelement, `node-name` to each named node, whose value consists of the node name. With this extension, the key for our example can be expressed as parts{}._{node-name, id}.

This does not alter any of the properties we expect to hold for keys and appears to account for any practical use of tag names in keys.

Acknowledgements. We are grateful to Byron Choi, Hartmut Liefke, Arnaud Sahuguet, Keishi Tajima, Chris Brew and Henry Thompson for a number of useful comments and discussions.

References

[1] Serge Abiteboul, Peter Buneman, and Dan Suciu. *Data on the Web. From Relations to Semistructured Data and XML.* Morgan Kaufman, 2000.

[2] Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases.* Addison-Wesley, 1995.

[3] Vidur Apparao, Steve Byrne, Mike Champion, Scott Isaacs, Ian Jacobs, Arnaud Le Hors, Gavin Nicol, Jonathan Robie, Robert Sutor, Chris Wilson, and Lauren Wood. *Document Object Model (DOM) Level 1 Specification.* W3C Recommendation, October 1998. `http://www.w3.org/TR/REC-DOM-Level-1/`.

[4] Arnaud Sahuguet. Everything You Ever Wanted to Know About DTDs, But Were Afraid to Ask. In *WebDB-2000*, 2000.

[5] Tim Bray, Jean Paoli, and C. M. Sperberg-McQueen. *Extensible Markup Language (XML) 1.0.* World Wide Web Consortium (W3C), Feb 1998. `http://www.w3.org/TR/REC-xml`.

[6] Peter Buneman, Susan Davidson, Wenfei Fan, Carmem Hara, and Wang-Chiew Tan. Reasoning about keys for XML. Draft manuscript, 2000.

[7] Byron Choi and Arnaud Sahuguet. DTD Inquisitor Demonstration. `http://xml.cis.upenn.edu/DTDi/`.

[8] James Clark and Steve DeRose. *XML Path Language (XPath).* W3C Working Draft, November 1999. `http://www.w3.org/TR/xpath`.

[9] Andrew Layman, Edward Jung, Eve Maler, and Henry S. Thompson. *XML-Data.* W3C Note, January 1998. `http://www.w3.org/TR/1998/NOTE-XML-data`.

[10] Raghu Ramakrishnan and Johannes Gehrke. *Database Management Systems.* McGraw-Hill Higher Education, 2000.

[11] Henry S. Thompson, David Beech, Murray Maloney, and Noah Mendelsohn. *XML Schema Part 1: Structures.* W3C Working Draft, April 2000. `http://www.w3.org/TR/xmlschema-1/`.

[12] Philip Wadler. A Formal Semantics for Patterns in XSL. Technical report, Computing Sciences Research Center, Bell Labs, Lucent Technologies, 2000. `http://www.cs.bell-labs.com/~wadler/topics/xml#xsl-semantics`.