Identifying Query Incompatibilities with Evolving XML Schemas

Pierre Genevès, Nabil Layaïda, Vincent Quint

► To cite this version:

Pierre Genevès, Nabil Layaïda, Vincent Quint. Identifying Query Incompatibilities with Evolving XML Schemas. ICFP'09, Aug 2009, Edinburgh, United Kingdom. pp.221-230, 10.1145/1596550.1596583 . inria-00423058

HAL Id: inria-00423058

<https://hal.inria.fr/inria-00423058>

Submitted on 9 Oct 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Identifying Query Incompatibilities with Evolving XML Schemas

Pierre Genevès

CNRS

pierre.geneves@inria.fr

Nabil Layaïda

Vincent Quint

INRIA

{nabil.layaida,vincent.quint}@inria.fr

Abstract

During the life cycle of an XML application, both schemas and queries may change from one version to another. Schema evolutions may affect query results and potentially the validity of produced data. Nowadays, a challenge is to assess and accommodate the impact of these changes in evolving XML applications. Such questions arise naturally in XML static analyzers. These analyzers often rely on decision procedures such as inclusion between XML schemas, query containment and satisfiability. However, existing decision procedures cannot be used directly in this context. The reason is that they are unable to distinguish information related to the evolution from information corresponding to bugs. This paper proposes a predicate language within a logical framework that can be used to make this distinction.

We present a system for monitoring the effect of schema evolutions on the set of admissible documents and on the results of queries. The system is very powerful in analyzing various scenarios where the result of a query may not be anymore what was expected. Specifically, the system is based on a set of predicates which allow a fine-grained analysis for a wide range of forward and backward compatibility issues. Moreover, the system can produce counterexamples and witness documents which are useful for debugging purposes. The current implementation has been tested with realistic use cases, where it allows identifying queries that must be reformulated in order to produce the expected results across successive schema versions.

Categories and Subject Descriptors D.3.4 [Software]: Programming Languages—Processors; D.2.4 [Software]: Engineering—Software/Program Verification

General Terms Languages, Standardization, Verification

Keywords XML, Schema, Queries, Evolution, Compatibility

1. Introduction

XML is now commonplace on the web and in many information systems where it is used for representing all kinds of information resources, ranging from simple text documents such as RSS or Atom feeds to highly structured databases. In these dynamic environments, not only data are changing steadily but their schemas

also get modified to cope with the evolution of the real world entities they describe.

Schema changes raise the issue of data consistency. Existing documents and data that were valid with a certain version of a schema may become invalid on a new version of the schema (forward incompatibility). Conversely, new documents created with the latest version of a schema may be invalid on some previous versions (backward incompatibility).

In addition, schemas may be written in different languages, such as DTD, XML Schema, or Relax-NG, to name only the most popular ones. And it is common practice to describe the same structure, or new versions of a structure, in different schema languages. Document formats developed by W3C provide a variety of examples: XHTML 1.0 has both DTDs and XML Schemas, while XHTML 2.0 has a Relax-NG definition; the schema for SVG Tiny 1.1 is a DTD, while version 1.2 is written in Relax-NG; MathML 1.01 has a DTD, MathML 2.0 has both a DTD and an XML Schema, and MathML 3.0 is developed with a Relax-NG schema and is expected to have also a DTD and an XML Schema. An issue then is to make sure that schemas written in different languages are equivalent, *i.e.* they describe the same structure, possibly with some differences due to the expressivity of the language [Murata et al. 2005]. Another issue is to clearly identify the differences between two versions of the same schema expressed in different languages. Moreover, the issues of forward and backward compatibility of instances obviously remain when schema languages change from a version to another.

Validation, and then compatibility, is not the only purpose of a schema. Validation is usually the first step for safe processing of documents and data. It makes sure that documents and data are structured as expected and can then be processed safely. The next step is to actually access and select the various parts to be handled in each phase of an application. For this, query languages play a key role. As an example, when transforming a document with XSL, XPath queries are paramount to locate in the original document the data to be produced in the transformed document.

Queries are affected by schema evolutions. The structures they return may change depending on the version of the schema used by a document. When changing schema, a query may return nothing, or something different from what was expected, and obviously further processing based on this query is at risk.

These observations highlight the need for evaluating precisely and safely the impact of schema evolutions on existing and future instances of documents and data. They also show that it is important for software engineers to precisely know what parts of a processing chain have to be updated when schemas change. In this paper we focus on the XPath query language which is used in many situations while processing XML documents and data. The XSL transformation language was already mentioned, but XPath is also present in XLink and XQuery for instance.

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.

ICFP'09, August 31–September 2, 2009, Edinburgh, Scotland, UK.
Copyright © 2009 ACM 978-1-60558-332-7/09/08...\$10.00

2. Analysis Framework

The main contribution of this paper is a framework that allows the automatic verification of properties related to XML schema and query evolution. In particular, it offers the possibility of checking fine-grained properties of the behavior of queries with respect to successive versions of a given schema. The system can be used for checking whether schema evolutions require a particular query to be updated. Whenever schema evolutions may induce query malfunctions, the system is able to generate annotated XML documents that exemplify bugs, with the goal of helping the programmer to understand and properly overcome undesired effects of schema evolutions.

The system relies on a predicate language (presented in Section 4) specifically designed for studying schema and query compatibility issues when schemas evolve. In particular, predicates allow characterizing in a precise manner nodes subject to evolution. For instance, predicates allow to distinguish new nodes selected by the query after a schema change from new nodes that appear in the modified schema. Predicates also allow to describe nodes that appear in new regions of a schema compared to its original version, or even in a new context described by a particular XPath expression. Predicates, together with the composition language provided in the system allow to express and analyze complex settings.

The system has been fully implemented [Genevès and Layaïda 2009] and is outlined in Figure 1. It is composed of a parser for reading the text file description of the problem (which in turn use specific parsers for schemas, queries, logical formulas, and predicates), compilers for translating schemas and queries into their logical representations, a solver for checking satisfiability of logical formulas, and a counter example XML tree generator (described in [Genevès et al. 2008]).

We first introduce the data model we consider for XML documents, schemas and queries.

XML Trees with Attributes An XML document is considered as a finite tree of unbounded depth and arity, with two kinds of nodes respectively named elements and attributes. In such a tree, an element may have any number of children elements, and may carry zero, one or more attributes. Attributes are leaves. Elements are ordered whereas attributes are not, as illustrated on Figure 4. In this paper, we focus on the nested structure of elements and attributes, and ignore XML data values.

Type Constraints As an internal representation for tree grammars, we consider regular tree type expressions (in the manner of [Hosoya et al. 2005]), extended with constraints over attributes. Assuming a set of variables ranged over by x , we define a tree type expression as follows:

$\tau ::=$		tree type expression
$\emptyset$		empty set
$()$		empty sequence
$\tau \mid \tau$		disjunction
τ, τ		concatenation
$l(a) [\tau]$		element definition
x		variable
$\text{let } \overline{x} = \overline{\tau} \text{ in } \tau$		binder

The **let** construct allows binding one or more variables to associated formulas. Since several variables can be bound at a time, the notation $\overline{x} = \overline{\tau}$ is used for denoting a vector of variable bindings (possibly with mutual recursion).

We impose a usual restriction on the recursive use of variables: we allow unguarded (*i.e.* not enclosed by a label) recursive uses of variables, but restrict them to tail positions¹. With that restriction,

¹For instance, “**let** $x = l(a) [\tau], x \mid ()$ **in** x ” is allowed.

tree types expressions define regular tree languages. In addition, an element definition may involve simple attribute expressions that describe which attributes the defined element may (or may not) carry:

$a ::=$		attribute expression
$()$		empty list
$list \mid a$		disjunction
$list ::=$		attribute list
$list, list$		commutative concatenation
$l?$		optional attribute
l		required attribute
$\neg l$		prohibited attribute

We use the usual semantics of regular tree types found in [Hosoya et al. 2005] and [Genevès et al. 2008].

Our tree type expressions capture most of the schemas in use today [Murata et al. 2005]. In practice, our system provides parsers that convert DTDs, XML Schemas, and Relax NGs to this internal tree type representation. Users may thus define constraints over XML documents with the language of their choice, and, more importantly, they may refer to most existing schemas for use with the system.

Queries The set of XPath expressions we consider is given by the syntax shown on Figure 2. The semantics of XPath expressions is described in [Clark and DeRose 1999], and more formally in [Wadler 2000]. We observed that, in practice, many XPath expressions contain syntactic sugars that can also fit into this fragment. Figure 3 presents how our XPath parser rewrites some commonly found XPath patterns into the fragment of Figure 2, where the notation $(axis::nt)^k$ stands for the composition of k successive path steps of the same form: $\underbrace{axis::nt / \dots / axis::nt}_{k \text{ steps}}$.

$query ::=$		
$/path$		absolute path
$path$		relative path
$query \mid query$		union
$query \cap query$		intersection
$path ::=$		
$path/path$		path composition
$path[qualifier]$		qualified path
$axis::nt$		step
$qualifier ::=$		
$qualifier \text{ and } qualifier$		conjunction
$qualifier \text{ or } qualifier$		disjunction
$\text{not}(qualifier)$		negation
$path$		path
$path/@nt$		attribute path
$@nt$		attribute step
$nt ::=$		
σ		node test
$*$		node label
$axis ::=$		any node label
$\text{self} \mid \text{child} \mid \text{parent}$		tree navigation axis
$\text{descendant} \mid \text{ancestor}$		
$\text{descendant-or-self}$		
ancestor-or-self		
following-sibling		
preceding-sibling		
$\text{following} \mid \text{preceding}$		

Figure 2. XPath Expressions.

Figure 1. Framework Overview.

$$\begin{aligned}
nt[\text{position}() = 1] &\rightsquigarrow nt[\text{not}(\text{preceding-sibling}::nt)] \\
nt[\text{position}() = \text{last}()] &\rightsquigarrow nt[\text{not}(\text{following-sibling}::nt)] \\
nt[\text{position}() = \underbrace{k}_{k>1}] &\rightsquigarrow nt[(\text{preceding-sibling}::nt)^{k-1}] \\
\text{count}(\text{path}) = 0 &\rightsquigarrow \text{not}(\text{path}) \\
\text{count}(\text{path}) > 0 &\rightsquigarrow \text{path} \\
\text{count}(nt) > \underbrace{k}_{k>0} &\rightsquigarrow nt/(\text{following-sibling}::nt)^k
\end{aligned}$$

$$\begin{aligned}
&\text{preceding-sibling}::*[\text{position}() = \text{last}() \text{ and } \text{qualifier}] \\
&\rightsquigarrow \text{preceding-sibling}::*[\text{not}(\text{preceding-sibling}::*) \text{ and } \text{qualifier}]
\end{aligned}$$

Figure 3. Syntactic Sugars and their Rewritings.

The next Section presents the logic underlying the predicate language. Section 4 describes predicates for characterizing the impact of schema changes. Finally, experiments on realistic use cases are reported in Section 5.

3. Logical Setting

It is well-known that there exist bijective encodings between unranked trees (trees of unbounded arity) and binary trees. Owing to these encodings binary trees may be used instead of unranked trees without loss of generality. In the sequel, we rely on a simple “first-child & next-sibling” encoding of unranked trees. In this encoding, the first child of an element node is preserved in the binary tree representation, whereas siblings of this node are appended as right successors in the binary representation. Attributes are left unchanged by this encoding. For instance, Figure 5 presents how the sample tree of Figure 4 is mapped.

Figure 4. Sample XML Tree with Attributes.

The logic we introduce below, used as the core of our framework, operates on such binary trees with attributes.

3.1 Logical Formulas

The concrete syntax of logical formulas is shown on Figure 6, where the meta-syntax $\langle X \rangle^\oplus$ means one or more occurrences of

Figure 5. Binary Encoding of Tree of Figure 4.

X separated by commas. The reader can directly use this syntax for encoding formulas as text files to be used with the system [Genevès and Layaïda 2009]. This concrete syntax is used as a single unifying notation throughout all the paper.

$\varphi ::=$	formula
T	true
F	false
l	element name
p	atomic proposition
#	start context
$\varphi \mid \varphi$	disjunction
$\varphi \ \& \ \varphi$	conjunction
$\varphi \Rightarrow \varphi$	implication
$\varphi \Leftrightarrow \varphi$	equivalence
(φ)	parenthesized formula
$\neg \varphi$	negation
$\langle p \rangle \varphi$	existential modality
$\langle l \rangle T$	attribute named l
$\$X$	variable
$\text{let } \langle \$X = \varphi \rangle^\oplus \text{ in } \varphi$	binder for recursion
predicate	predicate (See Section 4)
$p ::=$	program inside modalities
1	first child
2	next sibling
-1	parent
-2	previous sibling

Figure 6. Concrete Syntax of Formulas.

The semantics of logical formulas corresponds to the classical semantics of a μ -calculus interpreted over finite tree structures. A formula is satisfiable iff there exists a finite binary tree with attributes for which the formula holds at some node. This is formally defined in [Genevès et al. 2007], and we review it informally below through a series of examples.

Sample Formula	Tree	XML
$a \ \& \ <1>b$		$\langle a \rangle \langle b \rangle \langle /a \rangle$
$a \ \& \ <1>(b \ \& \ <2>c)$		$\langle a \rangle \langle b \rangle \langle c \rangle \langle /a \rangle$
$e \ \& \ <-1>(d \ \& \ <2>g)$		$\langle d \rangle \langle e \rangle \langle /d \rangle \langle g \rangle$
$f \ \& \ <-2>(g \ \& \ \sim \langle 2 \rangle T)$	none	none

Table 1. Sample Formulas and Satisfying Trees.

There is a difference between an element name and an atomic proposition²: an element has one and only one element name, whereas it can satisfy multiple atomic propositions. We use atomic propositions to attach specific information to tree nodes, not related to their XML labeling. For example, the start context (a reserved atomic proposition) is used to mark the starting context nodes for evaluating XPath expressions.

The logic uses modalities for navigating in binary trees. A modality $\langle p \rangle \varphi$ can be read as follows: “there exists a successor node by program p such that φ holds at this successor”. As shown on Figure 6, a program p is simply one of the four basic programs $\{1, 2, -1, -2\}$. Program 1 allows navigating from a node down to its first successor, and program 2 allows navigating from a node down to its second successor. The logic also features converse programs -1 and -2 for navigating upward in binary trees, respectively from the first successor to its parent and from the second successor to its previous sibling. Table 1 gives some simple formulas using modalities for navigating in binary trees, together with sample satisfying trees, in binary and unranked tree representations.

The logic allows expressing recursion in trees through the recursive binder. For example the recursive formula:

$$\text{let } \$X = b \mid \langle 2 \rangle \$X \text{ in } \$X$$

means that either the current node is named b or there is a sibling of the current node which is named b . For this purpose, the variable $\$X$ is bound to the subformula $b \mid \langle 2 \rangle \$X$ which contains an occurrence of $\$X$ (therefore defining the recursion). The scope of this binding is the subformula that follows the “in” symbol of the formula, that is $\$X$. The entire formula can thus be seen as a compact recursive notation for an infinitely nested formula of the form:

$$b \mid \langle 2 \rangle (b \mid \langle 2 \rangle (b \mid \langle 2 \rangle (\dots)))$$

Recursion allows expressing global properties. For instance, the recursive formula:

$$\sim \text{let } \$X = a \mid \langle 1 \rangle \$X \mid \langle 2 \rangle \$X \text{ in } \$X$$

expresses the absence of nodes named a in the whole subtree of the current node (including the current node). Furthermore, the fixpoint operator makes possible to bind several variables at a time, which is specifically useful for expressing mutual recursion. For example, the mutually recursive formula:

```
let
  $X = (a & <2>$Y) | <1>$X | <2>$X,
  $Y = b | <2>$Y
in $X
```

asserts that there is a node somewhere in the subtree such that this node is named a and it has at least one sibling which is named b . Binding several variables at a time provides a very expressive yet succinct notation for expressing mutually recursive structural patterns (that are common in XML Schemas, for instance).

From a theoretical perspective, the recursive binder $\text{let } \$X = \varphi \text{ in } \psi$ corresponds to the fixpoint operators of the μ -calculus. It is shown in [Genevès et al. 2007] that the least fixpoint and the greatest fixpoint operators of the μ -calculus coincide over finite tree structures, for a restricted class of formulas called *cycle-free* formulas. Translations of XPath expressions and schemas presented in this paper always yield cycle-free formulas (see [Genevès et al. 2008] for more details).

3.2 Queries

The logic is expressive enough to capture the set of XPath expressions presented in Section 2. For example, Figure 7 illustrates how the sample XPath expression:

`child::r[child::w/@att]`

is expressed in the logic. From a given context in an XML document, this expression selects all r child nodes which have at least one w child with an attribute att . Figure 7 shows how it is expressed in the logic, on the binary tree representation. The formula holds for r nodes which are selected by the expression. The first part of the formula, φ , corresponds to the step `child::r` which selects candidates r nodes. The second part, ψ , navigates downward in the subtrees of these candidate nodes to verify that they have at least one immediate w child with an attribute att .

Translated Query: `child::r[child::w/@att]`

Translation:

$$\underbrace{r \ \& \ (\text{let } \$X = \langle -1 \rangle \# \mid \langle -2 \rangle \$X)}_{\varphi} \ \& \ \underbrace{\langle 1 \rangle \text{let } \$Y = w \ \& \ \langle att \rangle T \mid \langle 2 \rangle \$Y}_{\psi}$$

Figure 7. XPath Translation Example.

This example illustrates the need for converse programs inside modalities. The translated XPath expression only uses forward axes (child and attribute), nevertheless both forward and backward modalities are required for its logical translation. Without converse programs we would have been unable to differentiate selected nodes from nodes whose existence is simply tested. More generally, properties must often be stated on both the ancestors and the descendants of the selected node. Equipping the logic with both forward and converse programs is therefore crucial. Logics without converse programs may only be used for solving XPath emptiness but cannot be used for solving other decision problems such as containment efficiently.

A systematic translation of XPath expressions into the logic is given in [Genevès et al. 2007]. In this paper, we extended it to

² In practice, an atomic proposition must start with a “_”.

deal with attributes. We implemented a compiler that takes any expression of the fragment of Figure 2 and computes its logical translation. With the help of this compiler, we extend the syntax of logical formulas with a logical predicate `select("query",  $\varphi$ )`. This predicate compiles the XPath expression *query* given as parameter into the logic, starting from a context that satisfies φ . The XPath expression to be given as parameter must match the syntax of the XPath fragment shown on Figure 2 (or Figure 3). In a similar manner, we introduce the predicate `exists("query",  $\varphi$ )` which tests the existence of *query* from a context satisfying φ , in a qualifier-like manner (without moving to its result). Additionally, the predicate `select("query")` is introduced as a shortcut for `select("query", #)`, where # simply marks the initial context node of the XPath expression³. The predicate `exists("query")` is a shortcut for `exists("query", T)`. These syntactic extensions of the logic allow the user to easily embed XPath expressions and formulate decision problems out of them (like *e.g.* containment or any other boolean combination). In the next sections we explain how the framework allows combining queries with schema information for formulating problems.

3.3 Tree Types

Tree type expressions are compiled into the logic in two steps: the first stage translates them into binary tree type expressions, and the second step actually compiles this intermediate representation into the logic. The translation procedure from tree type expressions to binary tree type expressions is well-known and detailed in [Genevès 2006]. The syntax of output expressions follows:

$\tau ::=$	binary tree type expression
$\emptyset$	empty set
$()$	empty tree
$\tau \mid \tau$	disjunction
$l(a) [x, x]$	element definition
$\text{let } \overline{x_i} = \overline{\tau_i} \text{ in } \tau$	binder

Attribute expressions are not concerned by this transformation to binary form: they are simply attached, unchanged, to new (binary) element definitions. Finally, binary tree type expressions are compiled into the logic. This translation step was introduced and proven correct in [Genevès et al. 2007]. Originally, the translation takes a tree type expression τ and returns the corresponding logical formula. Here, we extend it slightly but crucially: the logical translation of an expression τ is given by the function $\text{tr}(\tau)_{\varphi}^{\psi}$ defined below, that takes additional arguments φ and ψ :

$$\begin{aligned} \text{tr}(\tau)_{\varphi}^{\psi} &\stackrel{\text{def}}{=} \mathbf{F} \quad \text{for } \tau = \emptyset, () \\ \text{tr}(\tau_1 \mid \tau_2)_{\varphi}^{\psi} &\stackrel{\text{def}}{=} \text{tr}(\tau_1)_{\varphi}^{\psi} \mid \text{tr}(\tau_2)_{\varphi}^{\psi} \\ \text{tr}(l(a) [x_1, x_2])_{\varphi}^{\psi} &\stackrel{\text{def}}{=} (l \ \& \ \varphi \ \& \ \text{tra}(a) \ \& \ s_1(x_1) \ \& \ s_2(x_2)) \mid \psi \\ \text{tr}(\text{let } \overline{x_i} = \overline{\tau_i} \text{ in } \tau)_{\varphi}^{\psi} &\stackrel{\text{def}}{=} \text{let } \$\overline{x_i} = \text{tr}(\tau_i)_{\varphi}^{\psi} \text{ in } \text{tr}(\tau)_{\varphi}^{\psi} \end{aligned}$$

The addition of φ and ψ (respectively in a new conjunction and a new disjunction) is a key element for the definition of predicates in Section 4. More precisely, this allows marking type sub-expressions so that they can be distinguished in predicates, as explained in Section 3.4. In addition, φ and ψ are either true, false, or simple atomic propositions. Thus, it is worth noticing that their addition does not affect the linear complexity of tree type translation. The function $s.(\cdot)$ describes the type for each successor:

$$s_p(x) = \begin{cases} \sim\langle p \rangle T & \text{if } x \text{ is bound to } () \\ \sim\langle p \rangle T \mid \langle p \rangle \$X & \text{if } \text{nullable}(x) \\ \langle p \rangle \$X & \text{if } \text{not nullable}(x) \end{cases}$$

³ This mark is especially useful for comparing two or more XPath expressions from the same context.

according to the predicate `nullable( $x$ )` which indicates whether the type $T \neq ()$ bound to x contains the empty tree.

The function `tra( $a$ )` compiles attribute expressions associated with element definitions as follows:

$$\begin{aligned} \text{tra}() &\stackrel{\text{def}}{=} \text{notothers}() \\ \text{tra}(list \mid a) &\stackrel{\text{def}}{=} \text{tra}(list) \ \& \ \text{notothers}(list) \\ \text{tra}(list, list') &\stackrel{\text{def}}{=} \text{tra}(list) \ \& \ \text{tra}(list') \\ \text{tra}(l?) &\stackrel{\text{def}}{=} l \mid \sim l \\ \text{tra}(l) &\stackrel{\text{def}}{=} l \\ \text{tra}(\sim l) &\stackrel{\text{def}}{=} \sim l \end{aligned}$$

In usual schemas (*e.g.* DTDs, XML Schemas) when no attribute is specified for a given element, it simply means no attribute is allowed for the defined element. This convention must be explicitly stated in the logic. This is the role of the function “`notothers( $list$ )`” which returns the negated disjunction of all attributes not present in *list*. As a result, taking attributes into account comes at an extra-cost. The above translation appends a (potentially very large) formula in which all attributes occur, for each element definition. In practice, a placeholder atomic proposition is inserted until the full set of attributes involved in the problem formulation is known. When the whole formula has been parsed, placeholders are replaced by the conjunction of negated attributes they denote. This extra-cost can be observed in practice, and the system allows two modes of operations: with or without attributes⁴. Nevertheless the system is still capable of handling real world DTDs (such as the DTD of XHTML 1.0 Strict) with attributes. This is due to (1) the limited expressive power of languages such as DTD that do not allow for disjunction over attribute expressions (like “`list | a`”); and, more importantly, (2) the satisfiability-testing algorithm which is implemented using symbolic techniques [Genevès et al. 2008].

Tree type expressions form the common internal representation for a variety of XML schema definition languages. In practice, the logical translation of a tree type expression τ are obtained directly from a variety of formalisms for defining schemas, including DTD, XML Schema, and Relax NG. For this purpose, the syntax of logical formulas is extended with a predicate `type("f",  $\cdot$ )`. The logical translation of an existing schema is returned by `type("f",  $l$ )` where *f* is a file path to the schema file and *l* is the element name to be considered as the entry point (root) of the given schema. Any occurrence of this predicate will parse the given schema, extract its internal tree type representation τ , compile it into the logic and return the logical formula $\text{tr}(\tau)_T^F$.

3.4 Type Tagging

A tag (or “color”) is introduced in the compilation of schemas with the purpose of marking all node types of a specific schema. A tag is simply a fresh atomic proposition passed as a parameter to the translation of a tree type expression. For example: $\text{tr}(\tau)_{\text{xhtml}}^F$ is the logical translation of τ where each element definition is annotated with the atomic proposition “xhtml”. With the help of tags, it becomes possible to refer to the element types in any context. For instance, one may formulate $\text{tr}(\tau)_{\text{xhtml}}^F \mid \text{tr}(\tau')_{\text{smil}}^F$ for denoting the union of all τ and τ' documents, while keeping a way to distinguish element types; even if some element names are shared by the two type expressions.

Tagging becomes even more useful for characterizing evolutions between successive versions of a single schema. In this setting, we need a way to distinguish nodes allowed by a newer

⁴ The optional argument “-attributes” must be supplied for attributes to be considered.

schema version from nodes allowed by an older version. This distinction must not be based only on element names, but also on content models. Assume for instance that τ' is a newer version of schema τ . If we are interested in the set of trees allowed by τ' but not allowed by τ then we may formulate:

$$\text{tr}(\tau')_{\text{T}}^{\text{F}} \& \sim \text{tr}(\tau)_{\text{T}}^{\text{F}}$$

If we now want to check more fine-grained properties, we may rather be interested in the following (tagged) formulation:

$$\text{tr}(\tau')_{\text{all}}^{\text{F}} \& \sim \text{tr}(\tau)_{\text{T}}^{\text{old.complement}}$$

In this manner, we can distinguish elements that were added in τ' and whose names did not occur in τ , from elements whose names already occurred in τ but whose content model changed in τ' , for instance. In practice, a type is tagged using the predicate $\text{type}("f", l, \varphi, \varphi')$ which parses the specified schema, converts it into its logical representation τ and returns the formula $\text{tr}(\tau)_{\varphi'}^{\text{F}}$. This kind of type tagging is useful for studying the consequences of schema updates over queries, as presented in the next sections.

4. Analysis Predicates

This section introduces the basic analysis tasks offered to XML application designers for assessing the impact of schema evolutions. In particular, we propose a mean for identifying the precise reasons for type mismatches or changes in query results under type constraints.

For this purpose, we build on our query and type expression compilers, and define additional predicates that facilitate the formulation of decision problems at a higher level of abstraction. Specifically, these predicates are introduced as logical macros with the goal of allowing system usage while focusing (only) on the XML-side properties, and keeping underlying logical issues transparent for the user. Ultimately, we regard the set of basic logical formulas (such as modalities and recursive binders) as an assembly language, to which predicates are translated.

We illustrate this principle with two simple predicates designed for checking backward-compatibility of schemas, and query satisfiability in the presence of a schema.

- The predicate $\text{backward\_incompatible}(\tau, \tau')$ takes two type expressions as parameters, and assumes τ' is an altered version of τ . This predicate is unsatisfiable iff all instances of τ' are also valid against τ . Any occurrence of this predicate in the input formula will automatically be compiled as $\text{tr}(\tau')_{\text{T}}^{\text{F}} \& \sim \text{tr}(\tau)_{\text{T}}^{\text{F}}$.
- The predicate $\text{non\_empty}(\text{"query"}, \tau)$ takes an XPath expression (with the syntax defined on Figure 2) and a type expression as parameters, and is unsatisfiable iff the query always returns an empty set of nodes when evaluated on an XML document valid against τ . This predicate compiles into $\text{select}(\text{"query"}, \text{tr}(\tau)_{\text{T}}^{\text{F}} \& \#)$ where the top-level predicate $\text{select}(\text{"query"}, \varphi)$ compiles the XPath expression *query* into the logic, starting from a context that satisfies φ , as explained in Section 3.2. This can be used to check whether the modification of the schema does not contradict any part of the query.

Notice that the predicate $\text{non\_empty}(\text{"query"}, \tau)$ can be used for checking whether a query that is valid⁵ against a schema remains valid with an updated version of a schema. In other terms, this predicate allows determining whether a query that must always return a non-empty result (whatever the tree on which it is evaluated) keeps verifying the same property with a new version of a schema.

⁵ We say that a query is *valid* iff its negation is unsatisfiable.

A second, more-elaborate, class of predicates allows formulating problems that combine both a query *query* and two type expressions τ, τ' (where τ' is assumed to be a evolved version of τ):

- $\text{new\_element\_name}(\text{"query"}, \tau, \tau')$ is satisfied iff the query *query* selects elements whose names did not occur at all in τ . This is especially useful for queries whose last navigation step contains a "*"^{F} node test and may thus select unexpected elements. This predicate is compiled into:

$$\sim \text{element}(\tau) \& \text{select}(\text{"query"}, \text{tr}(\tau')_{\text{T}}^{\text{F}})$$

where $\text{element}(\tau)$ is another predicate that builds the disjunction of all element names occurring in τ . In a similar manner, the predicate $\text{attribute}(\varphi)$ builds the logical disjunction of all attribute names used in φ .

- $\text{new\_region}(\text{"query"}, \tau, \tau')$ is satisfied iff the query *query* selects elements whose names already occurred in τ , but such that these nodes now occur in a new context in τ' . In this setting, the path from the root of the document to a node selected by the XPath expression *query* contains a node whose type is defined in τ' but not in τ as illustrated below:

The predicate $\text{new\_region}(\text{"query"}, \tau, \tau')$ is logically defined as follows:

$$\begin{aligned} \text{new\_region}(\text{"query"}, \tau, \tau') &\stackrel{\text{def}}{=} \\ &\text{select}(\text{"query"}, \text{tr}(\tau')_{\text{all}}^{\text{F}} \& \sim \text{tr}(\tau)_{\text{T}}^{\text{old.complement}}) \\ &\quad \& \sim \text{added\_element}(\tau, \tau') \\ &\quad \& \text{ancestor}(\text{old.complement}) \\ &\quad \& \sim \text{descendant}(\text{old.complement}) \\ &\quad \& \sim \text{following}(\text{old.complement}) \\ &\quad \& \sim \text{preceding}(\text{old.complement}) \end{aligned}$$

The previous definition heavily relies on the partition of tree nodes defined by XPath axes, as illustrated by Figure 8. The definition of $\text{new\_region}(\text{"query"}, \tau, \tau')$ uses an auxiliary predicate $\text{added\_element}(\tau, \tau')$ that builds the disjunction of all element names defined in τ' but not in τ (or in other terms, elements that were added in τ'). In a similar manner, the predicate $\text{added\_attribute}(\varphi, \varphi')$ builds the disjunction of all attribute names defined in τ' but not in τ .

The predicate $\text{new\_region}(\text{"query"}, \tau, \tau')$ is useful for checking whether a query selects a different set of nodes with τ' than with τ because selected elements may occur in new regions of the document due to changes brought by τ' .

- $\text{new\_content}(\text{"query"}, \tau, \tau')$ is satisfied iff the query *query* selects elements whose names were already defined in τ , but whose content model has changed due to evolutions brought by τ' , as illustrated below:

Figure 8. XPath axes: partition of tree nodes.

The definition of $\text{new\_content}(\text{"query"}, \tau, \tau')$ follows:

$$\begin{aligned} \text{new\_content}(\text{"query"}, \tau, \tau') &\stackrel{\text{def}}{=} \\ &\text{select}(\text{"query"}, \text{tr}(\tau')_{\text{all}}^F \ \& \ \sim \text{tr}(\tau)_{\text{T}}^{\sim \text{old\_complement}}) \\ &\quad \& \sim \text{added\_element}(\tau, \tau') \\ &\quad \& \sim \text{ancestor}(\text{added\_element}(\tau, \tau')) \\ &\quad \& \text{descendant}(\sim \text{old\_complement}) \\ &\quad \& \sim \text{following}(\sim \text{old\_complement}) \\ &\quad \& \sim \text{preceding}(\sim \text{old\_complement}) \end{aligned}$$

The predicate $\text{new\_content}(\text{"query"}, \tau, \tau')$ can be used for ensuring that XPath expressions will not return nodes with a possibly new content model that may cause problems. For instance, this allows checking whether an XPath expression whose resulting node set is converted to a string value (as in, e.g. XPath expressions used in XSLT “value-of” instructions) is affected by the changes from τ to τ' .

The previously defined predicates can be used to help the programmer identify precisely how type constraint evolutions affect queries. They can even be combined with usual logical connectives to formulate even more sophisticated problems. For example, let us define the predicate $\text{exclude}(\varphi)$ which is satisfiable iff there is no node that satisfies φ in the whole tree. This predicate can be used for excluding specific element names or even nodes selected by a given XPath expression. It is defined as follows:

$$\begin{aligned} \text{exclude}(\varphi) &\stackrel{\text{def}}{=} \\ &\sim \text{ancestor-or-self}(\text{descendant-or-self}(\varphi)) \end{aligned}$$

This predicate can also be used for checking properties in an iterative manner, refining the property to be tested at each step. It can also be used for verifying fine-grained properties. For instance, one may check whether τ' defines the same set of trees as τ modulo new element names that were added in τ' with the following formulation:

$$\sim(\tau \leq \tau') \ \& \ \text{exclude}(\text{added\_element}(\tau, \tau'))$$

This allows identifying that, during the type evolution from τ to τ' , the query results change has not been caused by the type extension but by new compositions of nodes from the older type.

In practice, instead of taking internal tree type representations (as defined in Section 2) as parameters, most predicates do actually take any logical formula as parameter, or even schema paths as parameters. We believe this facilitates predicates usage and, most notably, how they can be composed together. Figure 9 gives the syntax of built-in predicates as they are implemented in the system, where f is a file path to a DTD (.dtd), XML Schema (.xsd), or Relax NG (.rng). In addition of aforementioned predicates, the predicate

```
predicate ::=
  select("query")
  select("query",  $\varphi$ )
  exists("query")
  exists("query",  $\varphi$ )

  type("f", l)
  type("f", l,  $\varphi$ ,  $\varphi'$ )
  forward_incompatible( $\varphi$ ,  $\varphi'$ )
  backward_incompatible( $\varphi$ ,  $\varphi'$ )

  element( $\varphi$ )
  attribute( $\varphi$ )
  descendant( $\varphi$ )
  exclude( $\varphi$ )
  added_element( $\varphi$ ,  $\varphi'$ )
  added_attribute( $\varphi$ ,  $\varphi'$ )

  non_empty("query",  $\varphi$ )
  new_element_name("query", "f", "f'", l)
  new_region("query", "f", "f'", l)
  new_content("query", "f", "f'", l)
  predicate-name( $\langle \varphi \rangle^\oplus$ )
```

Figure 9. Syntax of Predicates for XML Reasoning.

$\text{descendant}(\varphi)$ forces the existence of a node satisfying φ in the subtree, and $\text{predicate-name}(\langle \varphi \rangle^\oplus)$ is a call to a custom predicate, as explained in the next section.

4.1 Custom Predicates

Following the spirit of predicates presented in the previous section, users may also define their own custom predicates. The full syntax of XML logical specifications to be used with the system is defined on Figure 10, where the meta-syntax $\langle X \rangle^\oplus$ means one or more occurrence of X separated by commas. A global problem specification can be any formula (as defined on Figure 6), or a list of custom predicate definitions separated by semicolons and followed by a formula. A custom predicate may have parameters that are instantiated with actual formulas when the custom predicate is called (as shown on Figure 9). A formula bound to a custom predicate may include calls to other predicates, but not to the currently defined predicate (recursive definitions must be made through the let binder shown on Figure 6).

<i>spec</i> ::=	φ	formula (see Fig. 6)
	<i>def</i> ; φ	
<i>def</i> ::=	$\text{predicate-name}(\langle l \rangle^{\oplus}) = \varphi'$	custom definition
	<i>def</i> ; <i>def</i>	list of definitions

Figure 10. Global Syntax for Specifying Problems.

Schema	Variables	Elements	Attributes
XHTML 1.0 basic DTD	71	52	57
XHTML 1.1 basic DTD	89	67	83
MathML 1.01 DTD	137	127	72
MathML 2.0 DTD	194	181	97

Table 2. Sizes of (Some) Considered Schemas.

5. Framework in Action

We have implemented the whole software architecture described in Section 2 and illustrated on Figure 1. The tool [Genevès and Layaïda 2009] is available online from:

<http://wam.inrialpes.fr/xml>

We have carried out extensive experiments of the system with real world schemas such as XHTML, MathML, SVG, SMIL (Table 2 gives details related to their respective sizes) and queries found in transformations such MathML content to presentation [Pietriga 2005]. We present two of them that show how the tool can be used to analyze different situations where schemas and queries evolve.

Evolution of XHTML Basic

The first test consists in analyzing the relationship (forward and backward compatibility) between XHTML basic 1.0 and XHTML basic 1.1 schemas. In particular, backward compatibility can be checked by the following command:

```
backward_incompatible("xhtml-basic10.dtd",
                     "xhtml-basic11.dtd", "html")
```

The test immediately yields a counter example as the new schema contains new element names. The counter example (shown below) contains a `style` element occurring as a child of `head`, which is not permitted in XHTML basic 1.0:

```
<html>
  <head>
    <title/>
    <style type="_otherV"/>
  </head>
  <body/>
</html>
```

The next step consists in focusing on the relationship between both schemas excluding these new elements. This can be formulated by the following command:

```
backward_incompatible("xhtml-basic10.dtd",
                     "xhtml-basic11.dtd", "html")
& exclude(added_element(
  type("xhtml-basic10.dtd", "html"),
  type("xhtml-basic11.dtd", "html")))
```

The result of the test shows a counter example document that proves that XHTML basic 1.1 is not backward compatible with XHTML basic 1.0 even if new elements are not considered. In

particular, the content model of the `label` element cannot have an `a` element in XHTML basic 1.0 while it can in XHTML basic 1.1. The counter example produced by the solver is shown below:

```
<html>
  <head>
    <object>
      <label>
        <a>
          <img/>
        </a>
      </label>
    </object>
    <param/>
  </head>
  <body/>
</html>
```

XML basic 1.0 validity error: element "a" is not declared in "label" list of possible children

Notice that we observed similar forward and backward compatibility issues with several other W3C normative schemas (in particular for the different versions of SMIL and SVG). Such backward incompatibilities suggests that applications cannot simply ignore new elements from newer schemas, as the combination of older elements may evolve significantly from one version to another.

MathML Content to Presentation Conversion

MathML is an XML format for describing mathematical notations and capturing both its structure and graphical structure, also known as Content MathML and Presentation MathML respectively. The structure of a given equation is kept separate from the presentation and the rendering part can be generated from the structure description. This operation is usually carried out using an XSLT transformation that achieves the conversion. In this test series, we focus on the analysis of the queries contained in such a transformation sheet and evaluate the impact of the schema change from MathML 1.0 to MathML 2.0 on these queries.

Most of the queries contained in the transformation represent only a few patterns very similar up to element names. The following three patterns are the most frequently used:

```
Q1: //apply[*[1][self::eq]]
Q2: //apply[*[1][self::apply]/inverse]
Q3: //sin[preceding-sibling::*[position()=last()
    and (self::compose or self::inverse)]]
```

The first test is formulated by the following command:

```
new_region("Q1", "mathml1.dtd", "mathml2.dtd", "math")
```

The result of the test shows a counter example document that proves that the query may select nodes in new contexts in MathML 2.0 compared to MathML 1.0. In particular, the query Q1 selects `apply` elements whose ancestors can be declare elements, as indicated on the document produced by the solver:

```
<math xmlns:solver="http://wam.inrialpes.fr/xml"
      solver:context="true">
  <declare>
    <apply solver:target="true">
      <eq/>
    </apply>
  </condition>
```

```
</declare>
</math>
```

Notice that the solver automatically annotates a pair of nodes related by the query: when the query is evaluated from a node marked with the attribute `solver:context`, the node marked with `solver:target` is selected. To evaluate the effect of this change, the counter example is filled with content and passed as an input parameter to the transformation. This shows immediately a bug in the transformation as the resulting document is not a MathML 2.0 presentation document. Based on this analysis, we know that the XSLT template associated with the match pattern Q1 must be updated to cope with MathML evolution from version 1.0 to version 2.0.

The next test consists in evaluating the impact of the MathML type evolution for the query Q2 while excluding all new elements added in MathML 2.0 from the test. This identifies whether old elements of MathML 1.0 can be composed in MathML 2.0 in a different manner. This can be performed with the following command:

```
new_content("Q2", "mathml.dtd", "mathml2.dtd", "math")
& exclude(added_element(type("mathml.dtd", "math"),
    type("mathml2.dtd", "math")))
```

The test result shows an example document that effectively combines MathML 1.0 elements in a way that was not allowed in MathML 1.0 but permitted in MathML 2.0.

```
<math xmlns:solver="http://wam.inrialpes.fr/xml"
    solver:context="true">
  <apply solver:target="true">
    <apply>
      <inverse/>
    </apply>
    <annotation-xml>
      <math/>
    </annotation-xml>
    <condition/>
  </apply>
</math>
```

Similarly, the last test consists in evaluating the impact of the MathML type evolution for the query Q3, excluding all new elements added in MathML 2.0 and counter example documents containing `declare` elements (to avoid trivial counter examples):

```
new_regions("Q3", "mathml.dtd", "mathml2.dtd", "math")
& exclude(added_element(type("mathml.dtd", "math"),
    type("mathml2.dtd", "math")))
& exclude(declare)
```

The counter example document shown below illustrates a case where the `sin` element occurs in a new context.

```
<math xmlns:solver="http://wam.inrialpes.fr/xml"
    solver:context="true">
  <apply>
    <annotation-xml>
      <math>
        <apply>
          <inverse/>
          <sin solver:target="true"/>
        </apply>
      </math>
    </annotation-xml>
  </apply>
</math>
```

Applying the transformation on previous examples yields documents which are neither MathML 1.0 nor MathML 2.0 valid. As

a result, the stylesheet cannot be used safely over documents of the new type without modifications. In addition, the required changes to the stylesheet are not limited to the addition of new templates for MathML 2.0 elements. The templates that deal with the composition of MathML 1.0 elements should be revised as well.

All the previous tests were processed in less than 30 seconds on an ordinary laptop computer running Mac OS X. The 30s correspond to the most complex use cases. Most complex means analyzing recursive forward/backward and qualified queries such as Q3, under evolution of large and heavily recursive schemas such as XHTML and MathML (large number of type variables, elements and attributes: see Table 2). These are the hardest cases measured in practice with the implementation. Most of other schemas and queries usually found in applications are much simpler than the ones presented in this paper and will obviously be solved much faster. Given the variety of schemas occurring in practice, we focused on the most complex W3C standard schemas. The accompanying full online implementation [Genevès and Layaida 2009] allows to run all the tests described in the paper as well as user-supplied ones. It shows intermediate compilation stages, generated formulae (in particular the translation of schemas into the logic), and reports on the performance of each step of the analysis.

6. Related Work

Schema evolution is an important topic and has been extensively explored in the context of relational, object-oriented, and XML databases. Most of the previous work for XML query reformulation is approached through reductions to relational problems [Beyer et al. 2005]. This is because schema evolution was considered as a storage problem where the priority consists in ensuring data consistency across multiple relational schema versions. In such settings, two distinct schemas and an explicit description of the mapping between them are assumed as input. The problem then consists in reformulating a query expressed in terms of one schema into a semantically equivalent query in terms of the other schema: see [Yu and Popa 2005] and more recently [Moon et al. 2008] with references thereof.

In addition to the fundamental differences between XML and the relational data model, in the more general case of XML processing, schemas constantly evolve in a distributed, independent, and unpredictable environment. The relations between different schemas are not only unknown but hard to track. In this context, one priority is to help maintaining query consistency during these evolutions, which is still considered as a challenging problem [Sedlar 2005, Rose 2004]. The absence of evolution analysis tools for XML/XPath contrasts with the abundance of tools and methods routinely used in relational databases.

The work found in [Moro et al. 2007] discusses the impact of evolving XML schemas on query reformulation. Based on a taxonomy of XML schema changes during their evolution, the authors provide informal – not exact nor systematic – guidelines for writing queries which are less sensitive to schema evolution. In fact, studying query reformulation requires at least the ability to analyze the relationship between queries. For this reason, a closely related work is the problem of determining query containment and satisfiability under type constraints [Benedikt et al. 2005, Colazzo et al. 2006, Genevès et al. 2007].

The work found in [Benedikt et al. 2005] studies the complexity of XPath emptiness and containment for various fragments (see [Benedikt and Koch 2006] and references thereof for a survey). In [Colazzo et al. 2004, 2006], a technique is presented for statically ensuring correctness of paths. The approach deals with emptiness of XPath expressions without reverse axes. The work presented in [Genevès et al. 2007] solves the more general problem of containment, including reverse axes.

The main distinctive idea pursued in this paper is to develop a logical approach for guiding schema and query evolution. In contrast to the previous use of logics for proving properties such as query emptiness or equivalence, the goal here is different in that we seek to provide the necessary tools to produce relevant knowledge when such relations do not hold. From a complexity point-of-view, it is worth noticing that the addition of predicates does not increase complexity for the underlying logic shown in [Genevès et al. 2007].

We would also like to emphasize that, to the best of our knowledge, this work is the first to provide precise analyses of XML evolution, that was tested on real life use cases (such as XHTML and MathML types) and complex queries (involving recursive and backward navigation). As a consequence, in this context, analysis tools such as type-checkers [Hosoya and Pierce 2003, Benzaken et al. 2003, Möller and Schwartzbach 2005, Gapeyev et al. 2006, Castagna and Nguyen 2008] do not match the expressiveness, typing precision, and analysis capabilities of the work presented here.

7. Conclusion

In this article, we present an application of a logical framework for verifying forward/backward compatibility issues caused by schemas and queries evolution. We provide evidence that such a framework can be successfully used to overcome the obstacles of the analysis of XML type and query evolution. This kind of analyses is widely considered as a challenging problem in XML programming. As mentioned earlier, the difficulty is twofold: first it requires dealing with large and complex language constructions such as XML types and queries, and second, it requires modeling and reasoning about evolution of such constructions.

The presented tool allows XML designers to identify queries that need reformulation in order to produce the expected results across successive schema versions. With this tool designers can examine precisely the impact of schema changes over queries, therefore facilitating their reformulation. We gave illustrations of how to use the tool for both schema and query evolution on realistic examples. In particular, we considered typical situations in applications involving W3C schemas evolution such as XHTML and MathML. The tool can be very useful for standard schema writers and maintainers in order to assist them enforce some level of quality assurance on compatibility between versions.

There are a number of interesting extensions to the proposed system. In particular, the set of predicates can be easily enriched to detect more precisely the impact on queries. For example, one can extend the tagging to identify separately every navigation step and qualifier in a query expression. This will help greatly in the identification and reformulation of the navigation steps or qualifiers affected by schemas evolution.

References

- Michael Benedikt and Christoph Koch. XPath leashed. submitted, 2006.
- Michael Benedikt, Wenfei Fan, and Floris Geerts. XPath satisfiability in the presence of DTDs. In *PODS '05*, pages 25–36. ACM Press, 2005. ISBN 1-59593-062-0. doi: <http://doi.acm.org/10.1145/1065167.1065172>.
- Véronique Benzaken, Giuseppe Castagna, and Alain Frisch. CDuce: An XML-centric general-purpose language. In *ICFP '03: Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming*, pages 51–63, New York, NY, USA, 2003. ACM Press. ISBN 1-58113-756-7.
- Kevin Beyer, Fatma Özcan, Sundar Saiprasad, and Bert Van der Linden. DB2/XML: designing for evolution. In *SIGMOD '05*, pages 948–952. ACM, 2005. ISBN 1-59593-060-4. doi: <http://doi.acm.org/10.1145/1066157.1066299>.
- Giuseppe Castagna and Kim Nguyen. Typed iterators for XML. In *ICFP*, pages 15–26, 2008.
- James Clark and Steve DeRose. XML path language (XPath) version 1.0, W3C recommendation, November 1999. <http://www.w3.org/TR/1999/REC-xpath-19991116>.
- Dario Colazzo, Giorgio Ghelli, Paolo Manghi, and Carlo Sartiani. Types for path correctness of XML queries. In *ICFP '04: Proceedings of the ninth ACM SIGPLAN international conference on Functional programming*, pages 126–137, New York, NY, USA, 2004. ACM Press. ISBN 1-58113-905-5.
- Dario Colazzo, Giorgio Ghelli, Paolo Manghi, and Carlo Sartiani. Static analysis for path correctness of XML queries. *J. Funct. Program.*, 16(4-5):621–661, 2006. ISSN 0956-7968.
- Vladimir Gapeyev, François Garillot, and Benjamin C. Pierce. Statically typed document transformation: An Xtatic experience. In *PLAN-X 2006: Proceedings of the International Workshop on Programming Language Technologies for XML*, volume NS-05-6 of *BRICS Notes Series*, pages 2–13, Aarhus, Denmark, January 2006. BRICS.
- Pierre Genevès. *Logics for XML*. PhD thesis, Institut National Polytechnique de Grenoble, December 2006. <http://www.pierresoft.com/pierre.geneves/phd.htm>.
- Pierre Genevès and Nabil Layaïda. The XML reasoning solver project, February 2009. <http://wam.inrialpes.fr/xml>.
- Pierre Genevès, Nabil Layaïda, and Alan Schmitt. Efficient static analysis of XML paths and types. In *PLDI '07*, pages 342–351. ACM Press, 2007. ISBN 978-1-59593-633-2. doi: <http://doi.acm.org/10.1145/1250734.1250773>.
- Pierre Genevès, Nabil Layaïda, and Alan Schmitt. Efficient static analysis of XML paths and types. Long version of [Genevès et al. 2007], Research Report 6590, INRIA, July 2008. URL <http://hal.inria.fr/inria-00305302/en/>.
- Haruo Hosoya and Benjamin C. Pierce. XDuce: A statically typed XML processing language. *ACM Trans. Inter. Tech.*, 3(2):117–148, 2003. ISSN 1533-5399.
- Haruo Hosoya, Jérôme Vouillon, and Benjamin C. Pierce. Regular expression types for XML. *ACM TOPLAS*, 27(1):46–90, 2005. ISSN 0164-0925. doi: <http://doi.acm.org/10.1145/1053468.1053470>.
- Anders Möller and Michael I. Schwartzbach. The design space of type checkers for XML transformation languages. In *Proc. Tenth International Conference on Database Theory, ICDT '05*, volume 3363 of *LNCs*, pages 17–36, London, UK, January 2005. Springer-Verlag.
- Hyun J. Moon, Carlo A. Curino, Alin Deutsch, and Chien-Yi Hou. Managing and querying transaction-time databases under schema evolution. In *VLDB '08*, pages 882–895. VLDB Endowment, 2008.
- Mirella M. Moro, Susan Malaika, and Lipyeow Lim. Preserving xml queries during schema evolution. In *WWW '07*, pages 1341–1342. ACM, 2007. ISBN 978-1-59593-654-7. doi: <http://doi.acm.org/10.1145/1242572.1242841>.
- Makoto Murata, Dongwon Lee, Murali Mani, and Kohsuke Kawaguchi. Taxonomy of XML schema languages using formal language theory. *ACM TOIT*, 5(4):660–704, 2005. ISSN 1533-5399. doi: <http://doi.acm.org/10.1145/1111627.1111631>.
- Emmanuel Pietriga. MathML content2presentation transformation, May 2005. <http://www.lri.fr/~pietriga/mathmlc2p/mathmlc2p.html>.
- Kristoffer H. Rose. The XML world view. In *DocEng '04: Proceedings of the 2004 ACM symposium on Document engineering*, pages 34–34, New York, NY, USA, 2004. ACM. ISBN 1-58113-938-1. doi: <http://doi.acm.org/10.1145/1030397.1030403>. URL <http://www.research.ibm.com/XML/Rose-DocEng2004.pdf>.
- Eric Sedlar. Managing structure in bits & pieces: the killer use case for XML. In *SIGMOD '05*, pages 818–821. ACM, 2005. ISBN 1-59593-060-4. doi: <http://doi.acm.org/10.1145/1066157.1066256>.
- Philip Wadler. Two semantics for XPath. Internal Technical Note of the W3C XSL Working Group, <http://homepages.inf.ed.ac.uk/wadler/papers/xpath-semantics/xpath-semantics.pdf>, January 2000.
- Cong Yu and Lucian Popa. Semantic adaptation of schema mappings when schemas evolve. In *VLDB '05*, pages 1006–1017. VLDB Endowment, 2005. ISBN 1-59593-154-6.