


INSTITUT FÜR INFORMATIK
Lehr- und Forschungseinheit für
Programmier- und Modellierungssprachen
Oettingenstraße 67, D-80538 München

—————
Ludwig—————
Maximilians —
Universität —
München ———



Towards a Declarative Query and Transformation Language for XML and Semistructured Data: Simulation Unification

François Bry and Sebastian Schaffert

Technical Report, Computer Science Institute, Munich, Germany
<http://www.pms.informatik.uni-muenchen.de/publikationen>
Forschungsbericht/Research Report PMS-FB-2002-2, February 2002

Towards a Declarative Query and Transformation Language for XML and Semistructured Data: Simulation Unification

François Bry and Sebastian Schaffert

Institute for Computer Science, University of Munich
<http://www.pms.informatik.uni-muenchen.de>

Abstract. The growing importance of XML as a data interchange standard demands languages for data querying and transformation. Since the mid 90es, several such languages have been proposed that are inspired from functional languages (such as XSLT [1]) and/or database query languages (such as XQuery [2]). This paper addresses applying logic programming concepts and techniques to designing a declarative, rule-based query and transformation language for XML and semistructured data.

The paper first introduces issues specific to XML and semistructured data such as the necessity of flexible “query terms” and of “construct terms”. Then, it is argued that logic programming concepts are particularly appropriate for a declarative query and transformation language for XML and semistructured data. Finally, a new form of unification, called “simulation unification”, is proposed for answering “query terms”, and it is illustrated on examples.

1 Introduction

This article addresses applying logic programming to XML and semistructured data querying and transformation. This issue is of growing importance in both, practice and research. XML now is the data interchange standard of choice in application areas such as e-commerce, molecular biology, and astronomy. Furthermore, with the XML application XHTML, XML is becoming the preferred standard for textual web contents. Also recent languages for static and/or animated graphics such as SVG (Scalable Vector Graphics), X3D (an XML-based redefinition of the Virtual Reality Modelling Language VRML), and SMIL (Synchronised Multimedia Integration Language) are XML applications. “Native XML” database management systems are already marketed.

As a consequence, the design and implementation of selector and query languages for XML such as CSS selectors and XPath are premier concerns of the World Wide Web Consortium (W3C). Selector languages such as CSS selectors and XPath have been developed initially for style-sheet and/or transformation languages such as CSS and XSLT. Selector languages are “path-oriented”, i.e. a node (i.e. a subterm) in a tree (i.e. a term) is specified in terms of a root-to-node path in the manner of the file selection formalisms of operating systems. Constructs from regular expressions languages such as *, +, |, ? and “wildcards” give rise to expressing node repetitions, options, and nodes with unspecified labels. These constructs are essential in selecting and/or querying XML and semistructured data, for a basic principle of XML and semistructured data is that data items, called “documents”, do not have to conform to a predefined

schema (expressed in XML in the DTD or XML Schema formalisms). This principle aims at facilitating the interchange of data in unnormed contexts like the World Wide Web. XML-like tree structured data items that do not necessarily conform to a schema are called “semistructured” in database research [3].

Query languages [4] inspired from SQL and OQL [5] have been developed for XML and semistructured data. Queries in these languages for XML and semistructured data in general consist of two parts: a query proper and a construct part. The query part expresses a selection of nodes (i.e. subterms) from a data item (i.e. term). Node selection is in general expressed in a path-oriented formalism à la XPath extended with “and” and “or” connectives. The construct parts serve to re-assemble into new terms (some of) the nodes (i.e. terms) selected in the query part. The construct part is also called restructuring or transformation part. For example, the following query from [6] (there with `books-with-prices` instead of `book-price-comparison`) expresses in XML Query [2], also called XQuery, the request to list for each book found as `entry` element at `amazon.com` and as `book` element at `bn.com` the book’s title and the prices at both sources:

```

<book-price-comparison>
  { for $a in document("www.amazon.com/reviews.xml")//entry,
    $b in document("www.bn.com/bib.xml")//book
    where $b/title = $a/title
    return
      <book-with-prices>
        { $b/title }
        <price-amazon>{ $a/price/text() } </price-amazon>
        <price-bn>{ $b/price/text() } </price-bn>
      </book-with-prices>
    }
</book-price-comparison>

```

The query part is contained between the first `{` and `return`. The construct part is specified in the `book-with-prices` element. Note the node selection (expressed with the XPath expressions `$a/price/text()` and `$b/price/text()`) contained in the construct part.

The work reported about in this paper is based upon the conviction that logic programming provides with concepts giving rise to query and transformation languages more declarative than those based on a path-oriented node selection. A query term inspired from a Prolog goal atom would give rise to a “context-conscious” selection of several nodes (i.e. subterms) within a same term (the “context”) at a time, a term inspired from a Prolog head atom would be a convenient construct expression, rules relating (conjunctions or disjunctions of) query terms to construct terms would define views (in the database sense) and give rise to a deduction-like (backward or forward) chaining of term constructions – a feature often needed in XML and semistructured data processing.

In such a language inspired from logic programming, the previous query example can be expressed as follows (symbols beginning with upper case letters denote variables):

```

construct
  <book-price-comparison>
    all <book-with-prices>
      <title> T</title>
      <price-amazon> Pa</price-amazon>
      <price-bn> Pb</price-bn>
    </book-with-prices>
  </book-price-comparison>
where
  in amazon.com:
    <entry>
      <title>T</title>
      <price>Pa</price>
    </entry>
  and
  in bn.com:
    <book>
      <title>T</title>
      <price>Pb</price>
    </book>

```

In a more conventional syntax, this query can be expressed as follows (an element name is shortened to the first letters of its constituting words and the locations `amazon.com` and `bn.com` are omitted):

$$\text{bpc}\{ \text{all bwp}[t[T], \text{pa}[Pa], \text{pb}[Pb]] \} \leftarrow e[t[T], p[Pa]] \text{ and } b[t[T], p[Pb]]$$

An advantage of such a rule is to clearly separate node selection, expressed only in the query terms i.e. in the rule body, from construction, expressed in the construct term i.e. in the rule head. This is beneficial for both, the programmer and query evaluation. Another advantage of the approach is to avoid the rather procedural navigation through data item imposed by a path-oriented node selection. In the rule given above, the contents of both elements `t` (i.e. `title`) and `pa` (i.e. `price-amazon`) are selected in a single query term `e[t[T], p[Pa]]`. In contrast, the XQuery expression needs two paths for the same selection, `$a/title` and `$a/price/`. The query term `e[t[T], p[Pa]]` stresses the common context and the relative position of the selected nodes (i.e. subterms) `T` and `Pa`. In contrast, the paths `$a/title` and `$a/price/` specify two independent navigations through a term. Arguably, a term-oriented (or context-conscious or positional) node selection is more declarative than a path-oriented (or navigational) node selection.

This paper reports about first achievements in designing a term-oriented, “context-conscious”, or “positional” query and transformation language for XML and semistructured data. In order to conform to the semistructured data paradigm, a novel form of unification is needed. This paper is mostly devoted to motivating and specifying a nonstandard unification, called “simulation unification” convenient for a positional querying and transformation of XML and semistructured data.

This article is organised as follows. Section 1 is this introduction. Section 2 describes those aspects of the query and transformation language under development that are relevant to this paper. Simulation unification is addressed in Section 3. Section 4 is devoted to related work and a conclusion.

2 Elements of a Query and Transformation Language

This section introduces into those aspects of an experimental query and transformation language for XML and semistructured data, called Xcerpt, that are relevant to this paper. Aspects of XML, such as attributes and namespaces, that are irrelevant to this paper, are not explicitly addressed in the following. Two disjoint sets of symbols, the set \mathcal{L} of labels (or tags) and the set \mathcal{V} of variables are considered. Labels (variables, resp.) are denoted by words starting with a lower (upper, resp.) case letter. The following meta-variables (with or without indices and/or superscripts) are used:

- l denotes a label,
- X denotes a variable,
- t denotes a term (as defined below).

2.1 Database Terms

Database terms are an abstraction of XML documents. Following a common practice in XML query language and semistructured data research [3], a database is a set (or multiset) of database terms and the children of a document node may be either ordered (as in SGML and in standard XML), or unordered (as in the semistructured data model). In the following, a term whose root is labelled l and has *ordered* children t_1, \dots, t_n is denoted $l[t_1, \dots, t_n]$; a term whose root is labelled l and has *unordered* children t_1, \dots, t_n is denoted $l\{t_1, \dots, t_n\}$.

Definition 1 (Database Terms). *Database terms are inductively defined as follows:*

1. *A label is a (atomic) database term.*
2. *If l is a label and t_1, \dots, t_n are $n \geq 1$ database terms, then $l[t_1, \dots, t_n]$ and $l\{t_1, \dots, t_n\}$ are database terms.*

Database terms are similar to classical logic ground terms except that, (1) the arity of a function symbol, called here “label”, is not fixed (as in Prolog), and (2) the arguments of a function symbol may be unordered.

Whatever storage is used, a database term $t_0 = l\{t_1, \dots, t_n\}$ with unordered subterms t_1, \dots, t_n will always be stored in a manner inducing an order on t_1, \dots, t_n . The notion of unordered subterms t_1, \dots, t_n means that (1) the storage ordering of t_1, \dots, t_n is left at the discretion of the storage system (giving rise e.g. to clustering as many t_i as possible on a secondary memory page), and (2) no given ordering is to be returned when t_0 is accessed.

In the following, \mathcal{T}_{db} denotes the set of all database terms.

2.2 Query Terms

A query term is a “pattern” that specifies a selection of database terms very much like Prolog goal atoms and SQL selections. However, answers to query terms (cf. below Definition 13) differ from answers to Prolog goal atoms and SQL selections as follows:

- Database terms with additional subterms to those explicitly mentioned in a query term might be answers to this query term.
- Database terms with a different subterm ordering from that of the query term might be answers to this query term.
- A query term might specify subterms at an arbitrary depth.

In query terms, the single square and curly brackets, $[]$ and $\{ \}$, denote “exact subterm patterns”, i.e. single (square or curly) brackets are used in a query term to be answered by database terms with no more subterms than those given in the query term. Double square and curly brackets, $[[]]$ and $\{\{ \}\}$, on the other hand, denote “partial subterm patterns” as described above.

$[]$ and $[[]]$ are used if the subterm order in the answers is to be that of the query term, $\{ \}$ and $\{\{ \}\}$ are used otherwise. Thus, possible answers to the query term $t_1 = a[b, c\{\{d, e\}\}, f]$ are the database terms $a[b, c\{d, e, g\}, f]$ and $a[b, c\{d, e, g\}, f\{g, h\}]$ and $a[b, c\{d, e\{g, h\}\}, g], f\{g, h\}]$ and $a[b, c\{d, e\}, f]$. In contrast, $a[b, c\{d, e\}, f, g]$ and $a\{b, c\{d, e\}, f\}$ are no answers to t_1 . The only answers to $f\{ \}$ are f -labelled database terms with no children.

The construct *descendant*, short *desc*, introduces a subterm at an unspecified depth. Thus, possible answers to the query term $t_2 = a[\text{desc } f[c, d], b]$ are $a[f[c, d], b]$ and $a[g[f[c, d]], b]$ and $a[g[f[c, d], h], b]$ and $a[g[g[f[c, d]]], b]$ and $a[g[g[f[c, d], h], i], b]$.

In a query term, a variable X can be restricted to some query terms using the construct \rightsquigarrow , read “as”. Thus, the query term $t_3 = a[X_1 \rightsquigarrow b[[c, d]], X_2, e]$ constrains the variable X_1 to such database terms that are possible answers to the query term $b[[c, d]]$. Note that the variable X_2 is unconstrained in t_3 . Possible answers to t_3 are e.g. $a[b[c, d], f, e]$ which binds X_1 to $b[c, d]$ and X_2 to f , $a[b[c, d], f[g, h], e]$ which binds X_1 to $b[c, d]$ and X_2 to $f[g, h]$, $a[b[c, d, e], f, e]$ which binds X_1 to $b[c, d, e]$ and X_2 to f , and $a[b[c, e, d], f, e]$ which binds X_1 to $b[c, e, d]$ and X_2 to f .

Definition 2 (Query Terms). *Query terms are inductively defined as follows:*

1. If l is a label, then l and $l\{ \}$ are (atomic) query terms.
2. A variable X is a query term.
3. If X is a variable and t a query term, then $X \rightsquigarrow t$ is a query term.
4. If X is a variable and t is a query term, then $X \rightsquigarrow \text{desc } t$ is a query term.
5. If l is a label and t_1, \dots, t_n are $n \geq 1$ query terms, then $l[t_1, \dots, t_n]$, $l\{t_1, \dots, t_n\}$, $l[[t_1, \dots, t_n]]$, and $l\{\{t_1, \dots, t_n\}\}$ are query terms.

Multiple variable constraints are not precluded. A possible answer to e.g. $a\{\{X \rightsquigarrow b\{c\}\}, X \rightsquigarrow b\{d\}\}$ is $a\{b\{c, d\}\}$. The query term $a[[X \rightsquigarrow b\{c\}\}, X \rightsquigarrow f\{\{d\}\}]$, however, has no answers, as the labels b and f are distinct.

Subterms (of query terms) are defined as usual (e.g. a and X and $Y \rightsquigarrow \text{desc } b\{X\}$ and $h\{a, X \rightsquigarrow k\{c\}\}$ and $X \rightsquigarrow k\{c\}$ and t itself are subterms of $t = f\{a, g\{Y \rightsquigarrow \text{desc } b\{X\}\}, h\{a, X \rightsquigarrow k\{c\}\}\}$). In the following, query terms are assumed to be variable well-formed, a notion defined as follows.

Definition 3 (Variable Well-Formed Query Terms). *A term variable X depends on a term variable Y in a query term t if $X \rightsquigarrow t_1$ is a subterm of t and Y is a subterm of t_1 . A query term t is variable well-formed if t contains no term variables X_0, \dots, X_n ($n \geq 1$) such that 1. $X_0 = X_n$ and 2. for all $i = 1, \dots, n$, X_i depends on X_{i-1} in t .*

E.g. $f\{X \rightsquigarrow g\{X\}\}$ and $f\{X \rightsquigarrow g\{Y\}, Y \rightsquigarrow h\{X\}\}$ are not variable well-formed. Variable well-formedness precludes queries specifying infinite answers. Usually terms that are not variable well-formed are called cyclic. However, Xcerpt also allows for arbitrary graph structures (which are not discussed in this paper, cf. [7]) which might be cyclic in another sense.

In the following, query terms are implicitly assumed to be variable well-formed and the set \mathcal{T}_q is defined as the set of all (variable well-formed) query terms.

2.3 Construct Terms

Construct terms serve to re-assemble variables, the “values” of which are specified in query terms, so as to form new database terms. Thus, construct terms may contain both constructs $[]$ and $\{\}$ (like database terms) as well as variables. However, the construct \rightsquigarrow is not allowed in construct terms, as variables should be constrained where they are defined, (i.e. in query terms), not in construct terms where they are used to specify new terms.

Definition 4 (Construct Terms). *Construct terms are inductively defined as follows:*

1. *A label l is a (atomic) construct term.*
2. *A variable X is a construct term.*
3. *If l is a label and t_1, \dots, t_n are $n \geq 1$ construct terms, then $l[t_1, \dots, t_n]$ and $l\{t_1, \dots, t_n\}$ are construct terms.*

The set of construct terms will be denoted with \mathcal{T}_c in the rest of this paper. Note that $\mathcal{T}_{db} \subseteq \mathcal{T}_c \subseteq \mathcal{T}_q$.

2.4 Construct-Query Rules

Construct-query rules, short rules, relate queries, consisting of a conjunction of query terms, and construct terms. It is assumed (cf. below Point 3 of Definition 5) that each variable occurring in the construct term of a construct-query rule also occurs in at least one of the query terms of the rule, i.e. variables in construct-query rules are assumed to be “range-restricted” or “allowed”. A relaxation of this condition like in Prolog does not seem to be desirable.

Definition 5 (Construct-Query Rule). *A construct-query rule is an expression of the form $t^c \leftarrow t_1^q \wedge \dots \wedge t_n^q$ such that:*

1. *$n \geq 1$ and for all $i = 1, \dots, n$, t_i^q is a query term,*
2. *t^c is a construct term, and*

3. every variable occurring in t^c also occurs in at least one of the t_i^q .

The left hand-side, i.e. the construct term, of a (construct-query) rule will be referred to as the rule “head”. The right hand-side of a (construct-query) rule will be referred to as the rule “body”. Note that, in contrast to the body of a Prolog clause, the body of a (construct-query) rule cannot be empty, for empty rule bodies do not seem to be needed for the applications considered.

An Xcerpt program consists of a finite set of (construct-query) rules with a (conjunction of) query term(s). The scope of an occurrence of a variable in an Xcerpt program is, like in Prolog, restricted to the rule it occurs in.

2.5 Further Features

The preceding sections describe only part of the features of the experimental language Xcerpt. Additional features, although not relevant to the present paper, are outlined in the following, so as to give a more accurate impression of the Xcerpt project.

Database, query, and construct terms may, in addition to subterms, also have attributes à la XML, i.e. “flat” subterms subject to certain querying rules.

In database and construct terms, nestings of $[]$ and $\{ \}$ are possible, like in $a[b, \{c, d, e\}, f]$, with the obvious meaning that b (f , resp.) should precede (follow, resp.) the subterms with roots labelled c , d , and e . Similarly, nestings of $[]$, $\{ \}$, $[[]]$ and $\{\{ \}\}$, like in $a[[b, \{c, d, e\}, f]]$ and $a[b, \{\{c, d, e\}\}, f]$ are possible in query terms. In query terms, the descendant construct *desc* does not have to be preceded with a $X \rightsquigarrow$ expression, like in $f\{a, desc\ g\{X\}, c\}$.

In query terms, multiplicities can be specified like in the query term $a[a, 2-5\ b\{c, d\}, e]$. This query term admits as answers only those database terms having between 2 and 5 subterms that are answers to the query (sub)term $b\{c, d\}$, e.g. $a[a, b[c, d], b[c[f], d[f]], e]$. Also, label restrictions and specifications (e.g. in the form of regular expressions) and label variables are possible in query terms. Descendant specifications might be restricted with both, multiplicities and/or label restrictions or specifications. Subterm positions such as *first*, *last*, and numbered positions are also expressible, like in text processing languages.

In construct terms, the construct *all* (*some*, resp.) serves to collect in a database term all possible bindings (a nondeterministically chosen binding, resp.) of a variable. E.g. if each possible binding of the variable X is a publications of Anna, then the construct term $publications[author[Anna], \{all\ X\}]$ specifies a term containing all publications of Anna. In contrast, $publications[author[Anna], some\ X]$ specifies a term containing one randomly selected publication of Anna. Such a random selection is often needed in databases. It corresponds to Prolog cut (!) and *once* constructs. The *some* construct can be further specified with a multiplicity n or a multiplicity interval $n-m$.

Disjunctions of query atoms as well as arithmetic and string operations on labels may occur in bodies of construct-query rules. Also negations (to be processed as failure) of query atoms may occur in bodies of construct-query rules.

The syntax used in the present paper slightly departs from the Xcerpt original syntax. For example, $a[X \rightsquigarrow b\{c, d\}, e]$ is expressed in Xcerpt original syntax as:

```
element a [ X as element b element c ; element d ; element e ]
```

Following [8], Xcerpt programs are representable in XML, i.e. Xcerpt has an additional XML-based syntax to be used e.g. for data interchange.

Some issues deserve more investigations. The rich (and redundant) system of references of XML (with id/idref attributes and links) needs a representation in Xcerpt. A polymorphic type system giving rise to specifying abstract datatypes à la XML Schema [9,10] is desirable. The semantics of Xcerpt, especially the semantics of Xcerpt nonmonotonic negation, deserves further investigation. Compositionality and referential transparency are sought for.

3 Simulation Unification

The rule-based language Xcerpt, the main elements of which have been introduced above in Section 2, can be processed by both forward and backward chaining. Techniques similar to those used in implementations of Prolog (e.g. the use of the run-time stack for implementing a depth-first search) or of Datalog (e.g. a database storage of goal atoms) can be used for Xcerpt as well. However, Xcerpt cannot rely on standard unification because of the requirements on query terms listed in Section 2.2: A query term of the form $l[[t_1, \dots, t_n]]$ or $l\{\{t_1, \dots, t_n\}\}$ should “unify” with l -labelled terms with more subterms than those matching t_1, \dots , and t_n ; also unordered subterms (like in $l\{\{t_1, \dots, t_n\}\}$), the descendant construct *desc* and the *as* construct \rightsquigarrow have to be dealt with. This section is devoted to introducing a nonstandard unification called “simulation unification” fulfilling these requirements.

For space reasons, simulation unification is defined in this paper under the assumptions that $\{\{ \}$ and $\{ \}$ are the only kinds of braces, and that braces are only allowed immediately on the right of a label (like in $f\{\{a, g\{b, c\}, d\}\}$) and not directly within other braces (like in $f\{\{a, \{b, c\}, d\}\}$). In Section 3.5, we give a short overview on how to extend it to arbitrary term nestings involving $\{ \}$, $[\]$, $\{\{ \}$ and $[[\]]$. A formalisation integrating these constructs into the core algorithm is currently being worked on.

3.1 Simulation

Intuitively, a simulation of a graph G_1 in a graph G_2 is a mapping of the nodes of G_1 in the nodes of G_2 preserving the edges. In other words, there exists a simulation of G_1 in G_2 , if the node/edge structure of G_1 can be found as a subgraph of G_2 . Efficient algorithms for computing simulation (bisimulation, resp.) are given e.g. in [11]. In [3,12], simulation is used for verifying the conformity of semistructured data to a schema. The language UnQL [13] introduces (bi)simulation for query answering, but the usage is restricted to pattern matching.

Definition 6 (Graph Simulation). Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs and let \sim be an equivalence relation on $V_1 \cup V_2$. A relation $\mathcal{S} \subseteq V_1 \times V_2$ is a simulation with respect to \sim of G_1 in G_2 if:

1. If $v_1 \mathcal{S} v_2$, then $v_1 \sim v_2$.
2. If $v_1 \mathcal{S} v_2$ and $(v_1, v'_1) \in E_1$, then there exists $v'_2 \in V_2$ such that $v'_1 \mathcal{S} v'_2$ and $(v_2, v'_2) \in E_2$.

A simulation \mathcal{S} of a graph $G_1 = (V_1, E_1)$ in a graph $G_2 = (V_2, E_2)$ is total, if for each $v_1 \in V_1$ there exists at least one $v_2 \in V_2$ such that $v_1 \mathcal{S} v_2$.

A simulation \mathcal{S} of a tree T_1 with root r_1 in a tree T_2 with root r_2 is a rooted simulation of T_1 in T_2 if $r_1 \mathcal{S} r_2$.

Note that the definition of a simulation \mathcal{S} of G_1 in G_2 does not preclude that two distinct vertices v_1 and v'_1 of G_1 are simulated by the same vertex v_2 of G_2 , i.e. $(v_1, v_2) \in \mathcal{S}$ and $(v'_1, v_2) \in \mathcal{S}$. Figure 1 gives examples of simulations (represented by the dashed edges) with respect to vertex label equality.




Fig. 1. Rooted Simulations (with respect to label equality)

The following result immediately follows from Definition 6.

Lemma 1. *Every rooted simulation is total.*

It is obvious that $\mathcal{S} = \emptyset$ is always a simulation from a graph into another graph, albeit not an interesting one. For our purposes, the interesting simulations are the total simulation between trees which include the rooted simulations. Total simulations that are functions, also called minimal total simulations, are usually referred to as graph homomorphisms. Simulation with respect to label equality is a first notion towards a formalisation of answers to query terms: If a database term t^{db} is to be an answer to a query term t^q (both terms being considered as trees), then there must exist a rooted simulation with respect to label equality of (the term/tree with no \rightsquigarrow and *desc* constructs subjacent to) t^q in t^{db} .

3.2 Term Lattice

Definition 7 (Ground Query Term). *A query term is ground if it contains no variables, no \rightsquigarrow and no *desc*.*

In the following, the set of all ground query terms, extended by the two special terms \perp (the “empty” term) and \top (the “full” term) will be denoted by $\mathcal{T}_{\text{ground}}$. Note that $\mathcal{T}_{\text{ground}} \neq \mathcal{T}_{\text{db}}$, since in contrast to database terms ground query terms may contain both constructs $\{ \}$ and $\{ \{ \}$. In order to cope with these constructs, it is necessary to extend the notion of a simulation for ground query terms.

Definition 8 (Ground Query Term Simulation). *Let $t_1 \in \mathcal{T}_{\text{ground}}$ and $t_2 \in \mathcal{T}_{\text{ground}}$. Let $S_i \subseteq \mathcal{T}_{\text{ground}}$ denote the set of subtrees of t_i ($i \in \{1, 2\}$). A relation $\mathcal{S} \subseteq S_1 \times S_2$ is a simulation of t_1 in t_2 if:*

1. $t_1 \mathcal{S} t_2$
2. If $l_1 \mathcal{S} l_2$ then $l_1 = l_2$.
3. If $l_1 \{ \{ t_1^1, \dots, t_n^1 \} \} \mathcal{S} l_2 \{ \{ t_1^2, \dots, t_m^2 \} \}$, then $l_1 = l_2$ and for all $i \in \{1, \dots, n\}$ there exists $j \in \{1, \dots, m\}$ such that $t_i^1 \mathcal{S} t_j^2$
4. If $l_1 \{ \{ t_1^1, \dots, t_n^1 \} \} \mathcal{S} l_2 \{ t_1^2, \dots, t_m^2 \}$, then $l_1 = l_2$ and for all $i \in \{1, \dots, n\}$ there exists $j \in \{1, \dots, m\}$ such that $t_i^1 \mathcal{S} t_j^2$
5. If $l_1 \{ t_1^1, \dots, t_n^1 \} \mathcal{S} l_2 \{ \{ t_1^2, \dots, t_m^2 \} \}$, then $l_1 = l_2$ and for all $i \in \{1, \dots, n\}$ there exists $j \in \{1, \dots, m\}$ such that $t_i^1 \mathcal{S} t_j^2$, and for all $j \in \{1, \dots, m\}$ there exists $i \in \{1, \dots, n\}$ such that $t_i^1 \mathcal{S} t_j^2$
6. If $l_1 \{ t_1^1, \dots, t_n^1 \} \mathcal{S} l_2 \{ t_1^2, \dots, t_m^2 \}$, then $l_1 = l_2$ and for all $i \in \{1, \dots, n\}$ there exists $j \in \{1, \dots, m\}$ such that $t_i^1 \mathcal{S} t_j^2$, and for all $j \in \{1, \dots, m\}$ there exists $i \in \{1, \dots, n\}$ such that $t_i^1 \mathcal{S} t_j^2$

Definition 9 (Simulation Preorder). \preceq is the preorder on $\mathcal{T}_{\text{ground}} \setminus \{ \perp, \top \}$ defined by $t_1 \preceq t_2$ if there exists a ground query term simulation of t_1 in t_2 .

The preorder \preceq is not an order, for although $t_1 = f\{a\} \preceq t_2 = f\{a, a\}$ and $t_2 = f\{a, a\} \preceq t_1 = f\{a\}$ (both a of t_2 can be simulated by the same a of t_1), $t_1 = f\{a\} \neq t_2 = f\{a, a\}$.

However, \preceq induces as follows a (partial) order on $\mathcal{T}_{\text{ground}}$. First, consider the equivalence relation \equiv on $\mathcal{T}_{\text{ground}}$ defined by the bisimulation $t_1 \equiv t_2$ if both, $t_1 \preceq t_2$ and $t_2 \preceq t_1$ hold. Since \preceq is reflexive and transitive, \equiv is also reflexive and transitive. \equiv is by definition symmetric.

It is natural to chose as representative of an equivalence class of $\mathcal{T}_{\text{ground}} / \equiv$ the class member with the minimal number of repeated subterms, e.g. $f\{a\}$ is chosen as representative of class $\{f\{a\}, f\{a, a\}, f\{a, a, a\}, f\{a, a, a, a\}, \dots\} \in \mathcal{T}_{\text{ground}} / \equiv$.

In the following, referring to this representative will always be meant as a reference to the whole equivalence class and the (partial) order induced by \preceq on $\mathcal{T}_{\text{ground}} / \equiv$ will be noted \preceq , too. In other words, answers to query terms will be defined up to \equiv as representatives of elements of $\mathcal{T}_{\text{ground}} / \equiv$. Intuitively, $t_1 \preceq t_2$ means that it is possible to remove from t_2 subterms at arbitrary depth, until the remaining term is either t_1 or some \preceq -smaller term from the same \equiv -class as t_1 .

Definition 10 (Ground Query Term Lattice). \preceq is extended to \perp and \top as follows: For all $t \in \mathcal{T}_{\text{ground}}$, $\perp \preceq t$ and $t \preceq \top$. $(\mathcal{T}_{\text{ground}} / \equiv, \preceq)$ is the ground query term lattice.

3.3 Answers

An answer in a database $D \subseteq \mathcal{T}_{db}$ to a query term t^q is characterised by a set of values for the variables in t^q such that the ground query term t_g^q resulting from substituting these values for the variables in t^q is simulated by an element t of D (i.e. $t_g^q \preceq t$).

Consider for example the query $t^q = f\{\{X \rightsquigarrow g\{\{b\}\}, X \rightsquigarrow g\{\{c\}\}\}\}$ against the database $D = \{f\{g\{a, b, c\}, g\{a, b, c\}, h\}, f\{g\{b\}, g\{c\}\}\}$. The \rightsquigarrow constructs in t^q yield the constraint $g\{\{b\}\} \preceq X \wedge g\{\{c\}\} \preceq X$. The first database term in D yields the constraint $X \preceq g\{a, b, c\}$. The second database term in D yields the constraint $X \preceq g\{b\} \wedge X \preceq g\{c\}$. The constraint $g\{\{b\}\} \preceq X \wedge g\{\{c\}\} \preceq X$ is incompatible with $X \preceq g\{b\} \wedge X \preceq g\{c\}$. Thus, the only possible value for X is $g\{a, b, c\}$ and the only possible answer to t^q in D is $t_a^q = f\{g\{a, b, c\}, g\{a, b, c\}, h\}$.

Note that, in contrast to Prolog and SQL, the binding $X = g\{a, b, c\}$ does not suffice to characterise the answer t_a^q , for t^q does not have any “handle” for the subterm h of t_a^q . If not only the bindings for X but the complete answers to t^q are sought for, then the query term $Y \rightsquigarrow f\{\{X \rightsquigarrow g\{\{b\}\}, X \rightsquigarrow g\{\{c\}\}\}\}$ is to be used instead of t^q .

Definition 11 (Substitutions and Instances). *Let t^q be a query term and let X_1, \dots, X_n be the variables occurring (left or right of \rightsquigarrow or elsewhere) in t^q .*

A substitution is a function which assigns a construct term to each variable of a finite set of variables. A substitution σ is a grounding substitution for a query term t^q if σ assigns a ground query term to each variable in t^q .

If σ is a substitution (grounding substitution, resp.) for t^q assigning t_i to X_i ($1 \leq i \leq n$), then the instances (ground instances, resp.) of t^q with respect to σ are those construct terms (ground query terms, resp.) that can be constructed from t^q as follows:

1. *Replace each subterm $X \rightsquigarrow t$ by X .*
2. *Replace each occurrence of X_i by t_i ($1 \leq i \leq n$).*

Requiring in Definition 2 *desc* to occur to the right of \rightsquigarrow makes it possible to characterise ground instances of query terms by substitutions. This is helpful for formalising answers but not necessary for language implementations.

Not all ground instances of a query term are acceptable answers, for some instances might violate the conditions expressed by the \rightsquigarrow and *desc* constructs.

Definition 12 (Allowed Instances). *The constraint induced by a query term t^q and a substitution σ is the conjunction of all inequations $t\sigma \preceq X\sigma$ such that $X \rightsquigarrow t$ with $t \neq \text{desc } t_1$ is a subterm of t^q , and of all expressions $X\sigma \triangleleft t\sigma$ (read “ $X\sigma$ subterm of $t\sigma$ ”) such that $X \rightsquigarrow \text{desc } t$ is a subterm of t^q , if t^q has such subterms. If t^q has no such subterms, the constraint induced t^q and σ is the formula true.*

Let σ be a grounding substitution of a query term t^q . The instance $t\sigma$ of t^q is allowed if:

- $t\sigma \neq \perp$ and $t\sigma \neq \top$.

- Each inequality $t_1 \preceq t_2$ in the constraint induced by t^q and σ is satisfied in $(\mathcal{T} / \equiv, \preceq)$.
- If $t_1 \triangleleft t_2$ occurs in the constraint induced by t^q and σ , then there exists a subterm t'_1 of t_1 such that $t_2 \preceq t'_1$.

Definition 13 (Answers). Let t^q be a query term, D a database (i.e. $D \subseteq \mathcal{T}_{db}$). An answer to t^q in D is a database term $t^{db} \in D$ such that there exists an allowed instance t_a^q of t^q satisfying $t_a^q \preceq t^{db}$.

For some applications it might be preferable to restrict the answers to those database terms that are minimal in the sense that, at each level, they have no repetitions of identical subterms like in $f\{a, a\}$. This way, however, an answer would no longer necessarily be an element of the database considered. In all applications the authors are aware of, this is not desirable.

3.4 Simulation Unification

Simulation unification is a non-deterministic method for solving inequations of the form $t^q \preceq t^c$, where t^q is a query term, t^c is a construct term (possibly a database term), and t^q and t^c are variable disjoint, in the database term lattice $(\mathcal{T}_{db} / \equiv, \preceq)$, i.e. to determine substitutions σ such that $t^q\sigma$ and $t^c\sigma$ have instances $t^q\sigma\tau$ and $t^c\sigma\tau$ such that $t^q\sigma\tau$ and $t^c\sigma\tau$ are database terms and $t^q\sigma\tau \preceq t^c\sigma\tau$ holds.

Such inequations may result from a forward chaining evaluation of a query-construct rule against database terms. In such a case, the right-hand side t^c of the inequation contains no variables, i.e. it is a database term. An inequation $t^q \preceq t^c$ may also result from a backward chaining evaluation of the query term t^q against a query-construct rule whose head is t^c . In such a case, variables may occur in the construct term t^c but t^q and t^c are variable disjoint. That t^q and t^c do not share variables follows from the variable scoping rule for Xcerpt programs postulated in Section 2.4 above (this is the so-called “standardisation apart” of deduction methods).

Simulation unification consists in repeated applications of Term Decomposition phases followed by a Consistency Verification phase to a formula C (for constraint store) consisting in disjunctions of conjunctions of inequations of the form $t^q \preceq t^c$ (with t^q query term and t^c construct term) and/or equations of the form $t_1^c = t_2^c$ (with t_1^c and t_2^c construct terms). At the beginning C consists in a single inequation $t^q \preceq t^c$. Both phases Term Decomposition and Consistency Verification consist in stepwise changes of the constraint store C . These changes are expressed in the following formalism inspired from [14]:

A “simplification” $L \Leftrightarrow R$ replaces L by R .

Trivially satisfied inequations or equations are replaced by the atomic formula *true*. Inconsistent conjunctions of inequations or equations are replaced by the atomic formula *false*.

Definition 14 (Term Decomposition Rules). Let l (with or without indices) denote a label. Let t^1 and t^2 (with or without indices) denote query terms.

– **Root Elimination:**

$$\begin{aligned}
(1) \quad & l \preceq l\{t_1^2, \dots, t_m^2\} \Leftrightarrow \text{true} && \text{if } m \geq 1 \\
& l \preceq l\{\} \Leftrightarrow \text{true} \\
& l\{\} \preceq l\{t_1^2, \dots, t_m^2\} \Leftrightarrow \text{false} && \text{if } m \geq 1 \\
& l\{\} \preceq l \Leftrightarrow \text{true} \\
& l\{\} \preceq l\{\} \Leftrightarrow \text{true}
\end{aligned}$$

$$\begin{aligned}
(2) \quad & l\{t_1^1, \dots, t_n^1\} \preceq l \Leftrightarrow \text{false} && \text{if } n \geq 1 \\
& l\{t_1^1, \dots, t_n^1\} \preceq l\{\} \Leftrightarrow \text{false} && \text{if } n \geq 1 \\
& l\{t_1^1, \dots, t_n^1\} \preceq l \Leftrightarrow \text{false} && \text{if } n \geq 1 \\
& l\{t_1^1, \dots, t_n^1\} \preceq l\{\} \Leftrightarrow \text{false} && \text{if } n \geq 1
\end{aligned}$$

$$\begin{aligned}
(3) \quad & \text{Let } \Pi \text{ be the set of total functions } \{t_1^1, \dots, t_n^1\} \rightarrow \{t_1^2, \dots, t_m^2\}: \\
& l\{t_1^1, \dots, t_n^1\} \preceq l\{t_1^2, \dots, t_m^2\} \Leftrightarrow \bigvee_{\pi \in \Pi} \bigwedge_{1 \leq i \leq n} t_i^1 \preceq \pi(t_i^1) \\
& \hspace{15em} \text{if } n, m \geq 1
\end{aligned}$$

$$\begin{aligned}
& \text{Let } \Pi \text{ be the set of total, surjective functions } \{t_1^1, \dots, t_n^1\} \rightarrow \{t_1^2, \dots, t_m^2\}: \\
& l\{t_1^1, \dots, t_n^1\} \preceq l\{t_1^2, \dots, t_m^2\} \Leftrightarrow \bigvee_{\pi \in \Pi} \bigwedge_{1 \leq i \leq n} t_i^1 \preceq \pi(t_i^1) \\
& \hspace{15em} \text{if } n, m \geq 1
\end{aligned}$$

$$\begin{aligned}
(4) \quad & l_1\{t_1^1, \dots, t_n^1\} \preceq l_2\{t_1^2, \dots, t_m^2\} \Leftrightarrow \text{false if } l_1 \neq l_2 \text{ (} n, m \geq 0 \text{)} \\
& l_1\{t_1^1, \dots, t_n^1\} \preceq l_2\{t_1^2, \dots, t_m^2\} \Leftrightarrow \text{false if } l_1 \neq l_2 \text{ (} n, m \geq 0 \text{)}
\end{aligned}$$

– \rightsquigarrow **Elimination:**

$$X \rightsquigarrow t^1 \preceq t^2 \Leftrightarrow t^1 \preceq t^2 \wedge t^1 \preceq X \wedge X \preceq t^2$$

– **Descendant Elimination:**

$$\text{desc } t^1 \preceq l_2\{t_1^2, \dots, t_m^2\} \Leftrightarrow t^1 \preceq l_2\{t_1^2, \dots, t_m^2\} \vee \bigvee_{1 \leq i \leq m} \text{desc } t^1 \preceq t_i^2 \\
\hspace{15em} \text{if } m \geq 0$$

Applying the \rightsquigarrow and descendant elimination rules to a constraint store C in disjunctive normal form may yield a constraint store not in disjunctive normal form. Thus, the method has to restore from time to time the disjunctive normal form of C . In doing so, the formulas *true* and *false* are treated as usual: *true* is removed from conjunctions, conjunctions containing *false* are removed.

In the following, $mgcu(t_1, \dots, t_n)$ (with t_1, \dots, t_n construct terms) returns a most general commutative-unifier of t_1, \dots, t_n (in the sense of [15]) expressed as either *false*, if t_1 and t_2 are not commutative-unifiable, or as *true* if t_1 and t_2 are commutative-unifiable and do not contain variables, or else as a conjunction of equations of the form $X = t$. Note that most general commutative-unifiers are only computed for construct terms (i.e. terms without \rightsquigarrow and *desc* construct). Recall that commutative unification is decidable.

In the definition below, simulation unification is initialised with $X_0 \rightsquigarrow t^a \preceq t^c$, where X_0 is a variable occurring neither in t^a nor in t^c , instead of simply $t^a \preceq$

t^c . The additional variable X_0 serves to a complete specification of the answers returned. This is useful in proving the correctness of simulation unification but can usually be dispensed of in practice.

Definition 15 (Simulation Unification).

1. **Initialisation:** $C := X_0 \rightsquigarrow t^q \preceq t^c$
(with t^q query term, t^c construct term and t^q , t^c and X_0 variable disjoint).
2. **Term Decomposition:**
Until C can no longer be modified, repeat performing one of:
 - Apply a (applicable) Term Decomposition rule to C
 - Put C in disjunctive normal form
3. **Variable Binding:**
Replace each $X \preceq t$ in C with $X = t$.
4. **Consistency Verification:**
For each disjunct D of C and for each variable X occurring in D do:
Replace in D the equations $X = t_1, \dots, X = t_n$ by $\text{mgu}(t_1, \dots, t_n)$.

For efficiency reasons it is preferable to intertwine the Term Decomposition and Consistency Verification phases instead of performing them one after another. The sequential processing in Definition 15 simplifies the proofs.

Proposition 1 (Correctness and Completeness). *Let t^q be a query term, t^c a construct term, and X_0 a variable such that t^q , t^c and X_0 are variable disjoint. There exists a substitution τ such that $t^q\tau$ and $t^c\tau$ are database terms and $t^q\tau = t^c\tau$ if and only if a simulation unification initialised with $X_0 \rightsquigarrow t^q \preceq t^c$ returns a substitution σ such that*

- For each variable X in t^q , $X\sigma$ is a subterm of $t^q\sigma$.
- $t^q\tau$ is an instance of $t^q\sigma$.
- $t^c\tau$ is an instance of $t^c\sigma$.

The proof of Proposition 1 is given in Appendix.

3.5 Extension to the Term Nesting Constructs $[]$, $[[]]$, and $\{\{ \}\}$

Recall (cf. above Section 3.1) that rooted simulations couple the “nodes” of construct terms level-wise. Especially, if \mathcal{S} is a rooted simulation between two construct terms $l\{t_1^1, \dots, t_n^1\}$ and $l\{t_1^2, \dots, t_m^2\}$ (considered as trees) and if $(t_i^1, t) \in \mathcal{S}$ ($1 \leq i \leq n$), then t is one of the t_j^2 ($1 \leq j \leq m$).

In order to extend the previous definitions and results to arbitrary nestings, the simulation (cf. Definition 6) between two construct terms t^1 and t^2 to consider are those inducing between each level of t^1 and t^2 relations fulfilling the condition of Figure 2. The notion of monotonicity referred to in Figure 2 is defined as follows.

Definition 16 (Monotonic Relation). *Let $t^1 = l\{t_1^1, \dots, t_n^1\}$ and $t^2 = l\{t_1^2, \dots, t_m^2\}$ be construct terms. A relation $\mathcal{S} \subseteq \{t_1^1, \dots, t_n^1\} \times \{t_1^2, \dots, t_m^2\}$ is monotonic if for all $(t_i^1, t) \in \mathcal{S}$ ($1 \leq i \leq n$) $t \in \{t_i^2, t_{i+1}^2, \dots, t_m^2\}$.*

$t^1 \preceq t^2$	simulation / Π	Unification
$l\{\{t_1^1, \dots, t_n^1\} \preceq l\{t_1^1, \dots, t_m^1\}$		commutative unification
$l\{\{t_1^1, \dots, t_n^1\} \preceq l[t_1^1, \dots, t_m^1]$		commutative unification
$l[[t_1^1, \dots, t_n^1]] \preceq l[t_1^1, \dots, t_m^1]$	monotonic	simple unification
$l\{t_1^1, \dots, t_n^1\} \preceq l\{t_1^1, \dots, t_m^1\}$	surjective	commutative unification
$l\{\{t_1^1, \dots, t_n^1\} \preceq l[t_1^1, \dots, t_m^1]$	surjective	simple unification
$l[t_1^1, \dots, t_n^1] \preceq l[t_1^1, \dots, t_m^1]$	surjective and monotonic	simple unification

Fig. 2.

For processing the term nestings $\{ \}$, $[\]$, and $[[\]]$ during simulation unification it suffices to modify the cases (2) and (3) of rule Root Elimination (cf. Definition 14) as follows and to add to this rule the following cases (5) and (6):

- (2) $l\{t_1^1, \dots, t_n^1\} \preceq l \Leftrightarrow \text{false}$ if $n \geq 1$
 $l\{t_1^1, \dots, t_n^1\} \preceq l\{ \} \Leftrightarrow \text{false}$ if $n \geq 1$
... (as before)
 $l[t_1^1, \dots, t_n^1] \preceq l\{ \} \Leftrightarrow \text{false}$ if $n > 1$
 $l[[t_1^1, \dots, t_n^1]] \preceq l\{ \} \Leftrightarrow \text{false}$ if $n > 1$
 $l\{\{t_1^1, \dots, t_n^1\}\} \preceq l\{ \} \Leftrightarrow \text{false}$ if $n > 1$
- (3) Let t^1, t^2 and the set of (total) functions $\Pi : \{t_1^1, \dots, t_n^1\} \rightarrow \{t_1^2, \dots, t_m^2\}$ be as defined in Figure 2:
 $t^1 \preceq t^2 \Leftrightarrow \bigvee_{\pi \in \Pi} \bigwedge_{1 \leq i \leq n} t_i^1 \preceq \pi(t_i^1)$ if $n \geq 1$ and $m \geq 1$
- (5) $l[t_1^1, \dots, t_n^1] \preceq l\{t_1^1, \dots, t_m^1\} \Leftrightarrow \text{false}$
- (6) $l[[t_1^1, \dots, t_n^1]] \preceq l\{t_1^1, \dots, t_m^1\} \Leftrightarrow \text{false}$

Finally, the Consistency Verification phase of simulation unification must rely on a unification performing at each level of the terms compared as specified in Figure 2. In this table, “simple unification” denotes the standard unification as defined e.g. in [16].

3.6 Examples

$f\{\{X \rightsquigarrow b, Y \rightsquigarrow b\{\{c, d\}\}\}\}$ and $f\{a, b\{c, d, e\}, b\{e\}\}$ “simulation unify” yielding the following constraints: $(X = b\{c, d, e\} \wedge Y = b\{c, d, e\}) \vee (X = b\{e\} \wedge Y = b\{c, d, e\})$.

Also, the terms $X \rightsquigarrow \text{desc}(Y \rightsquigarrow f\{\{a\}\})$ and $g\{f\{Z, b, c\}, h\{f\{a, b\}\}\}$ “simulation unify” yielding $((Y = f\{Y, b, c\} \wedge a \preceq Z) \vee Y = f\{a, b\}) \wedge X = g\{f\{Z, b, c\}, h\{f\{a, b\}\}\}$.

Note that these simulation unifications constrain variables “on both sides”, i.e. simulation unification is no matching but a full-fledged unification.

4 Related Work and Conclusion

The articles [17,18,19] have already pointed out the drawbacks of relying on a navigational node selection à la XPath [20] and XQuery [2] for query and transformation languages for XML and semistructured data.

The language UnQL [13] has introduced simulation as a means for query answering. UnQL, like Xcerpt, uses the notions of patterns and templates. UnQL and Xcerpt differ from each other as follows. First, a query in UnQL consists of a single “select-where” expression which can be processed with pattern matching. In contrast, a query in Xcerpt might “chain” several “construct-query rules” requiring a “unification” which is capable of binding variables from both of the terms to be “unified”. Second, variables in UnQL can only occur as leaves of query patterns. Complex queries might require the use of several patterns in UnQL, where a single pattern suffices in Xcerpt.

In [17] a language for querying and transforming semistructured data is described. Like XPath and XQuery this language has variables for nodes, i.e. in the Xcerpt terminology labels.

[18] describes fxt, a language for querying and transforming semistructured data. fxt has variables for terms (or trees) and forests. fxt offers regular expressions similar to those of XPath for node selection. In contrast, the approach proposed in the present paper uses like Prolog variables for subterms. Arguably, languages with term variables makes data description less navigational than languages with node variables.

The language semantics in [17] is based upon a so-called component calculus and an algebra, very much in the style of XQuery’s algebra which is inspired from functional languages. The language semantics given in [18] for fxt is in terms of tree automata. Arguably, Definition 13 is closer to a Tarski’s style model theory and might therefore be seen as a more declarative semantics.

Several articles propose inference methods either rule-based or based upon consistency verification for XML data. [19] proposes a rule language very similar to Prolog called nowadays RuleML [21]. Several approaches that are too numerous for being explicitly mentioned here adapt techniques from feature logics to XML data. These approaches are usually named referring to “ontology” and/or “Semantic Web”. Common to RuleML and the ontology or Semantic Web approaches is that the language they propose do not support a direct access to XML data. Instead, their languages require a translation into a specific syntax. In some cases, like the binary predicate language RDF, this syntax might seem too stringent. For the authors of this paper, a direct access to XML data is an essential feature of an inference language for Web-based databases and semantic reasoning with Web data.

Simulation is no new notion. It is commonly used in process algebra and graph theory. It has been applied to semistructured data e.g. in [12,22,3] for schema validation. Graph simulation in general has been studied extensively cf. [11,23] (simulation is called “path inclusion” in [23]).

Several unification methods have been proposed that, like simulation unification, process flexible terms or structures, notably feature unification [24,25] and associative-commutative-unification, short AC-unification, [26]. Simulation

unification differs from feature unification in several aspects: First, it accepts terms with several identically “rooted” terms corresponding to a feature with more than one value. In contrast, feature logic languages and feature unification usually preclude this: A term (in Xcerpt syntax) such as $a\{f\{b\}, f\{c\}\}$ is not possible in a feature logic (although $a\{f\{b, c\}\}$ is possible). Second, ordered subterms, corresponding to ordered features, are not possible in feature logic languages and are not processed by feature unification. Third, simulation unification has a descendant construct which has no counterpart in (standard) feature logics (note, however, that a similar construct is considered in the work [27]). Finally, simulation unification of two Xcerpt terms is decidable while the satisfiability of general feature constraints with quantifiers is not decidable [25].

Simulation unification might remind of theory unification [28]. The significant difference between both is that simulation unification is based upon an order relation, while theory unification refers to a congruence relation.

There are interesting similarities between simulation unification and approaches to constraint solving over finite domains [29]. Simulation unification relies on a possibly disjunctive constraint store. This is rarely the case for constraint solvers. However, constraint programming approaches such as aggregation constraints [30] and constructive disjunction [31] seem interesting techniques for the future development of the language Xcerpt.

In this paper, a novel approach to querying and transforming XML and semistructured data based has been outlined. This approach is based on logic programming and a novel form of unification, simulation unification. A few aspects of a language under development, Xcerpt, have been presented. Many issues deserve further investigations. In particular, the complexity of simulation unification and its efficient implementation deserve further research.

Acknowledgements. The authors are thankful to Slim Abdennadher and Norbert Eisinger for useful suggestions.

References

1. W3C <http://www.w3.org/Style/XSL/>: Extensible Stylesheet Language (XSL). (2000)
2. W3C <http://www.w3.org/TR/xquery/>: XQuery: A Query Language for XML. (2001)
3. Abiteboul, S., Buneman, P., Suciu, D.: Data on the Web. From Relations to Semistructured Data and XML . Morgan Kaufmann Publishers, San Francisco, CA (2000)
4. Fernandez, M., Siméon, J., Wadler, P.: XML Query Languages: Experiences and Examples. Communication to the XML Query W3C Working Group (1999)
5. Alashqur, A.M., Su, S.Y.W., Lam., H.: OQL: A Query Language for Manipulating Object-Oriented Databases. In: Proc. 15th Int. Conf. on Very Large Data Bases (VLDB). (1989)
6. Chamberlin, D., Fankhauser, P., Marchiori, M., Robie, J.: XML Query Use Cases. W3C Working Draft 20 (2001)
7. Bry, F., Schaffert, S.: Pattern Queries for XML and Semistructured Data. Technical Report PMS-FB-2002-5, Inst. for Computer Sciences, University of Munich, <http://www.pms.informatik.uni-muenchen.de/publikationen/#PMS-FB-2002-5> (2002)
8. Maier, D.: Database Desiderata for an XML Query Language. In: Proceedings of QL’98 - The Query Languages Workshop. (1998) <http://www.w3.org/TandS/QL/QL98/>.
9. W3C <http://www.w3.org/TR/xmlschema-1/>: ML Schema Part 1: Structures. (2001)
10. W3C <http://www.w3.org/TR/xmlschema-2/>: ML Schema Part 2: Datatypes. (2001)
11. Henzinger, M.R., Henzinger, T.A., Kopke, P.W.: Computing Simulations on Finite and Infinite Graphs (1996)

12. Fernandez, M., Suciu, D.: Optimizing Regular Path Expressions Using Graph Schemas. In: Proceedings of the Int. Conf. on Data Engineering. (1988) 14–23
13. Buneman, P., Fernandez, M., Suciu, D.: UnQL: A Query Language and Algebra for Semistructured Data Based on Structural Recursion. VLDB Journal **9** (2000) 76–110
14. Frühwirth, T.: Theory and Practice of Constraint Handling Rules. Journal of Logic Programming, Special Issue on Constraint Logic Programming **37** (1998) 95–138
15. Baader, F.: Unification in Commutative Theories. In: Unification. Academic Press (1989) 417–435
16. Lloyd, J.W. In: Foundations of Logic Programming. Springer-Verlag (1993)
17. Grahne, G., Lakshmanan, L.V.S.: On the Difference between Navigating Semi-structured Data and Querying It. In: Workshop on Database Programming Languages. (1999)
18. Berlea, A., Seidl, H.: fxt – A Transformation Language for XML Documents. Journal of CIT, Special Issue on Domain-Specific Languages (2001)
19. Boley, H.: Relationships Between Logic Programming and XML. In: Proc. 14th Workshop Logische Programmierung, Würzburg (2000)
20. W3 Consortium <http://www.w3.org/TR/xpath>: XML Path Language (XPath). (1999)
21. DFKI: RuleML – Rule Markup Language. <http://www.dfki.uni-kl.de/ruleml/> (2002)
22. Buneman, P., Davidson, S.B., Fernandez, M.F., Suciu, D.: Adding Structure to Unstructured Data. In: Proceedings of ICDT'97. Volume 1186., Springer (1997) 336–350
23. Kilpeläinen, P.: Tree Matching Problems with Applications to Structured Text Databases. PhD thesis, Dept. of Computer Sciences, University of Helsinki (1992)
24. Aït-Kaci, H., Podelski, A., Goldstein, S.C.: Order-Sorted Theory Unification. Technical Report 32, digital – Paris Research Laboratory (1993)
25. Smolka, G.: Feature Constraint Logics for Unification Grammars. Journal of Logic Programming **12** (1992) 51–87
26. Fages, F.: Associative-Commutative Unification. In: Proc. 7th Int. Conf. on Automated Deduction (Napa, CA). Volume 170., Berlin, Springer (1984) 194–208
27. Backofen, R.: Regular Path Expressions in Feature Logic. Journal of Symbolic Computation **17** (1994) 412–455
28. Baader, F., Snyder, W.: Unification Theory. In Robinson, A., Voronkov, A., eds.: Handbook of Automated Reasoning. Elsevier Science Publishers (1999)
29. Montanari, U., Rossi, F.: Finite domain constraint solving and constraint logic programming. In Benhamou, F., Colmerauer, A., eds.: Constraint Logic Programming: Selected Research. MIT press (1993) 201–221
30. Ross, K.A., Srivastava, D., Stuckey, P.J., Sudarshan, S.: Foundations of aggregation constraints. Theoretical Computer Science B **190** (1994)
31. Würtz, J., Müller, T.: Constructive disjunction revisited. In: KI - Künstliche Intelligenz. (1996) 377–386

Appendix: Proof of Proposition 1

The proposition is established by reduction to the case where t^c is a database term (i.e. does not contain variables). With this reduction, the proposition can be reformulated as follows:

Lemma 2. *Let t^q be a query term, t^c be a database term. There exists a substitution σ such that $(X \rightsquigarrow t^q)\sigma = t^c$ if and only if a simulation unification initialized with $X \rightsquigarrow t^q \preceq t^c$ returns the substitution σ .*

This restriction can be lifted by defining a substitution μ such that $t^{db} = t^c\mu$ and $(X \rightsquigarrow t^q)\sigma = (X \rightsquigarrow t^q)\sigma\mu$. Since $(X \rightsquigarrow t^q)\sigma = t^{db} = t^c\mu$, and t^q and t^c are variable disjoint, follows $t^q\sigma\mu = t^c\sigma\mu$ and thus we can define τ as $\sigma\mu$.

We now continue with the proof of Lemma 2.

1 Sufficient Condition

It is to show that if the simulation unification returns a substitution σ , then $(X \rightsquigarrow t^q)\sigma = t^c$.

We prove the hypothesis through contradiction. Assume w.l.o.g. that the simulation unification would return a substitution σ that contains a single binding for a variable $Y = s$ in $(X \rightsquigarrow t^q)$ such that $(X \rightsquigarrow t^q)[s/Y] \neq t^c$.

Such a binding could only be the result of a \rightsquigarrow -elimination rule. Since the simulation unification is preserving the parent-child relationship in the root elimination, Y could only be bound to a subterm s of t^c at the same depth (of root eliminations), and thus the replacement of Y with s will yield just t^c , and thus the assumption does not hold, i.e. simulation unification is correct.

2 Necessary Condition

2.1 Completeness

It is to show that if there exists a σ such that $(X \rightsquigarrow t^q)\sigma = t^c$, then the simulation unification returns that σ .

The first rule application will be \rightsquigarrow -elimination, which will add especially the inequation $X \preceq t^c$. Since X is a variable that does not occur in t , there will be no second upper bound for X . In the variable binding algorithm, we will thus get the substitution $X = t^c$, which will trivially fulfill the condition $X \rightsquigarrow t\sigma = t^c$.

It thus remains to show that the result t^c is an allowed instance in the sense of definition 11. We have to show that if simulation unification returns σ , then $t\sigma \preceq t^c$. This is supported through the correctness shown below.

2.2 Termination

Lemma 3 (Termination). *For any query term t^q and construct term t^c such that t^q and t^c are variable disjoint, the simulation unification of $t^q \preceq t^c$ will terminate for any sequence of rule applications.*

Proof. For proving the termination of simulation unification, we make the following assumptions:

- the constraint store \mathcal{C} is initialized with the single inequation $t^q \preceq t^c$
- after each step, the DNF of the constraint store \mathcal{C} is re-created
- the transitivity rule is applied in a sensible manner, i.e. only if new information is added that has not yet been produced

We prove the proposition by induction. Each application of a rule will replace a disjunct with at least one (but finitely many) new disjuncts that all have a strictly lower *rank* than the old disjunct.

This is a sufficient requirement for the termination of the simulation unification: We can generate a tree for each possible sequence of rule applications where the root node is the formula containing the original inequation and the immediate successors are the disjuncts that result from applying one of the rules to a node. Each node has only finitely many (immediate) successors and each successor always has a strictly lower *rank* than the parent. Thus there cannot exist an infinite path and, according to *König's Lemma*, the tree is thus finite.

We define the *rank* of a formula of \wedge -connected inequations based on the *depth* of the participating terms (note that the *depth* is always ≥ 1).

- $depth(l) = 1$
- $depth(X) = 1$
- $depth(X \rightsquigarrow t) = 2 \cdot depth(t) + 2$
- $depth(l\{t_1, \dots, t_n\}) = depth(l\{\{t_1, \dots, t_n\}\}) = 1 + n \cdot \max_{i=1}^n (depth(t_i))$
- $depth(desc\ t) = 1 + depth(t)$
- $rank(t_1 \preceq t_2) = depth(t_1) \cdot depth(t_2)$
- $rank(S_1 \wedge S_2) = rank(S_1) + rank(S_2)$
- $rank(X = t) = 0$
- $rank(true) = rank(false) = 0$

Intuitively, the *depth* of a term equals the number of nodes of the *completion* of this term such that every node on the same level has the same degree and length to a leaf.

Root Elimination: Cases (1), (2) and (4) reduce to *true* or *false* and thus trivially satisfy the condition (*). In case (3), we have to show that

$$rank(f\{t_1^1, \dots, t_n^1\} \preceq f\{t_1^2, \dots, t_m^2\}) > rank(\bigwedge_{i=1}^n t_i^1 \preceq \pi(t_i^1))$$

On the left side, we get

$$\begin{aligned} rank(l\{t_1^1, \dots, t_n^1\} \preceq l\{t_1^2, \dots, t_m^2\}) &= \\ &= depth(l\{t_1^1, \dots, t_n^1\}) \cdot depth(l\{t_1^2, \dots, t_m^2\}) \\ &= (1 + n \cdot \max_{i=1}^n (depth(t_i^1))) \cdot (1 + m \cdot \max_{i=1}^m (depth(t_i^2))) \\ &= (1 + n \cdot \max_{i=1}^n (depth(t_i^1)) + m \cdot \max_{i=1}^m (depth(t_i^2))) \\ &\quad + n \cdot m \cdot \max_{i=1}^n (depth(t_i^1)) \cdot \max_{i=1}^m (depth(t_i^2)) \end{aligned}$$

For the right side, we get an upper estimation by

$$\begin{aligned}
\text{rank}\left(\bigwedge_{i=1}^n t_i^1 \preceq \pi(t_i^1)\right) &= \sum_{i=1}^n \text{depth}(t_i^1) \cdot \text{depth}(\pi(t_i^1)) \\
&\leq \sum_{i=1}^n \text{depth}(t_i^1) \cdot \max_{i=1}^m(\text{depth}(t_i^2)) \\
&\leq n \cdot \max_{i=1}^n(\text{depth}(t_i^1)) \cdot \max_{i=1}^m(\text{depth}(t_i^2)) \\
&\leq m \cdot n \cdot \max_{i=1}^n(\text{depth}(t_i^1)) \cdot \max_{i=1}^m(\text{depth}(t_i^2)) \\
&< m \cdot n \cdot \max_{i=1}^n(\text{depth}(t_i^1)) \cdot \max_{i=1}^m(\text{depth}(t_i^2)) + 1
\end{aligned}$$

Since this is part of the left side, case (3) also satisfy the condition (*).

The proof is analogous for the case with double curly brackets in the left term.

\rightsquigarrow **Elimination:** Again, we have to show that

$$\text{rank}(X \rightsquigarrow t^1 \preceq t^2) > \text{rank}(t^1 \preceq X \wedge X \preceq t^2 \wedge t^1 \preceq t^2)$$

The left side is evaluated as follows:

$$\begin{aligned}
\text{rank}(X \rightsquigarrow t^1 \preceq t^2) &= (2 \cdot \text{depth}(t^1) + 2) \cdot \text{depth}(t^2) \\
&= 2 \cdot \text{depth}(t^1) \cdot \text{depth}(t^2) + 2 \cdot \text{depth}(t^2)
\end{aligned}$$

Again, we calculate an upper estimation for the right side:

$$\begin{aligned}
\text{rank}(t^1 \preceq X \wedge X \preceq t^2 \wedge t^1 \preceq t^2) &= \\
&= \text{depth}(t^1) \cdot \text{depth}(X) + \text{depth}(X) \cdot \text{depth}(t^2) + \text{depth}(t^1) \cdot \text{depth}(t^2) \\
&= \text{depth}(t^1) + \text{depth}(t^2) + \text{depth}(t^1) \cdot \text{depth}(t^2) \\
&\leq \text{depth}(t^1) \cdot \text{depth}(t^2) + \text{depth}(t^2) + \text{depth}(t^1) \cdot \text{depth}(t^2) \\
&= 2 \cdot \text{depth}(t^1) \cdot \text{depth}(t^2) + \text{depth}(t^2) \\
&< 2 \cdot \text{depth}(t^1) \cdot \text{depth}(t^2) + 2 \cdot \text{depth}(t^2) \quad (\text{with } \text{depth}(t^2) \geq 1)
\end{aligned}$$

Descendant Elimination: We have to show that

- (1) $\text{rank}(\text{desc } t^1 \preceq l\{t_1^2, \dots, t_m^2\}) > \text{rank}(t^1 \preceq l\{t_1^2, \dots, t_m^2\})$
- (2) $\text{rank}(\text{desc } t^1 \preceq l\{t_1^2, \dots, t_m^2\}) > \text{rank}(\text{desc } t^1 \preceq t_i^2)$ for all $1 \leq i \leq m$

For the left side, we have in both cases:

$$\text{rank}(\text{desc } t^1 \preceq l\{t_1^2, \dots, t_m^2\}) = (1 + \text{depth}(t^1)) \cdot (1 + m \cdot \max_{i=1}^m(\text{depth}(t_i^2)))$$

The right side in case **1** is evaluated to

$$\text{rank}(t^1 \preceq l\{t_1^2, \dots, t_m^2\}) = \text{depth}(t^1) \cdot (1 + m \cdot \max_{i=1}^m(\text{depth}(t_i^2)))$$

which is trivially smaller than the left side. In case **2**, we get ($1 \leq i \leq m$):

$$\begin{aligned}
\text{rank}(\text{desc } t^1 \preceq t_i^2) &= (1 + \text{depth}(t^1)) \cdot \text{depth}(t_i^2) \\
&\leq (1 + \text{depth}(t^1)) \cdot \max_{i=1}^m(\text{depth}(t_i^2))
\end{aligned}$$

which again is trivially smaller than the left side.

We have thus shown that the application of the decomposition rules will terminate for any two terms t^q and t^c .

□