THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A query language and optimization techniques for unstructured data

Citation for published version:

Buneman, P, Davidson, S, Hillebrand, G & Suciu, D 1996, A query language and optimization techniques for unstructured data. in *SIGMOD '96 Proceedings of the 1996 ACM SIGMOD international conference on Management of data*. ACM Press, pp. 505-516. <https://doi.org/10.1145/233269.233368>

Digital Object Identifier (DOI):

[10.1145/233269.233368](https://doi.org/10.1145/233269.233368)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

SIGMOD '96 Proceedings of the 1996 ACM SIGMOD international conference on Management of data

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A Query Language and Optimization Techniques for Unstructured Data

Peter Buneman

University of Pennsylvania
peter@central.cis.upenn.edu

Susan Davidson*

University of Pennsylvania
susan@central.cis.upenn.edu

Gerd Hillebrand

University of Pennsylvania
gerdh@saul.cis.upenn.edu

Dan Suciu

AT&T Research
suciu@research.att.com

Abstract

A new kind of data model has recently emerged in which the database is not constrained by a conventional schema. Systems like ACeDB, which has become very popular with biologists, and the recent Tsimmis proposal for data integration organize data in tree-like structures whose components can be used equally well to represent sets and tuples. Such structures allow great flexibility in data representation

What query language is appropriate for such structures? Here we propose a simple language UnQL for querying data organized as a rooted, edge-labeled graph. In this model, relational data may be represented as fixed-depth trees, and on such trees UnQL is equivalent to the relational algebra. The novelty of UnQL consists in its programming constructs for arbitrarily deep data and for cyclic structures. While strictly more powerful than query languages with path expressions like XSQL, UnQL can still be efficiently evaluated. We describe new optimization techniques for the deep or “vertical” dimension of UnQL queries. Furthermore, we show that known optimization techniques for operators on flat relations apply to the “horizontal” dimension of UnQL.

1 Introduction

There are two good reasons to query and manipulate data whose structure is not constrained by a schema. First, some systems and proposals have recently emerged in which the schema is either absent or places very loose constraints on the data; second, for the purposes of browsing it may be convenient to forget the schema, even though one exists. In this new “unstructured” approach to data representation each com-

ponent, or object, is interpreted dynamically and may be linked to other components in an arbitrary fashion. We represent such data in a very general framework, which is roughly speaking an edge-labeled rooted directed graph. We call this model a “labeled tree” model because the query language we will develop is most easily thought of as a tree-traversing language, even though all expressible queries are well-defined on cyclic structures. To justify this choice of model let us briefly examine some current systems that can be conveniently represented in such a framework.

Unstructured data models. Biological data storage poses a problem for “fixed schema” systems, because the rapid evolution of experimental techniques requires constant adjustment of the schema [GRS93]. It is also convenient to have a model that accommodates missing data. One such database system that is extremely popular within the molecular biology community is ACeDB (A *C. elegans* Database) [TMD92]. Although ACeDB has a schema, it imposes only weak constraints on the database. A class in an ACeDB schema can be thought of as a labeled tree with non-terminal edges labeled by attribute names or base types and leaves labeled by base types or by other class names. Superficially, this is like an object-oriented model, and one could think of adapting an object-oriented query language. However, one would need to deal with the problem that the internal structure of each instance is a tree, and one would need to find a better way of dealing with missing data than that provided by object-oriented query languages. The query language for ACeDB allows selections of objects and pointer traversals, but it cannot perform general projections or joins.¹

Another system that uses a tree-like model is Tsimmis [PGMW95], which has been proposed for heterogeneous data integration. In Tsimmis there is no schema, and the “type” of data is interpreted by the user from labels in the structure. In particular, the difference between a record and a set in a Tsimmis structure is that

Corresponding author. Address: Department of Computer and Information Science, University of Pennsylvania, 200 South 33rd Street, Philadelphia, PA 19104-6389. Phone: (215) 898-3490. Fax: (215) 898-0587.

¹The recent TableMaker extension to ACeDB enables projections and joins, but it is a hack and does not form a compositional query language.

Figure 1: An example movie database.

within a record node the edge labels are all distinct, whereas within a set node the edge labels are all the same. A query language has been proposed for Tsimmis, but it allows only limited forms of “deep” traversal of the data structures. Languages for restructuring Tsimmis would also appear to be important, for Tsimmis is proposed as a model for data exchange. Moreover optimization of such languages appears to be an open issue.

Browsing. Even if a schema exists for the data (e.g., data stored in a relational database management system), it may be convenient to ignore it for browsing purposes. For example, one might want to search all the character strings in the database to determine which tables contain a particular term. One cannot write a generic relational algebra expression to express such a query. Languages such as XSQL [KKS92] provide a solution to this, and these can be extended, with some limitations, to object-oriented databases. However the problem of “deep” queries and query optimization remains. An example of a “deep” query is known to most programmers: find a file which is somewhere in a tree-structured file system. Utilities to perform this particular query are usually available, but they do not generalize easily, and dealing with recursive structures (which can arise from the presence of symbolic links) is usually achieved through a hack.

Edge-labeled trees. As an example of the structure we want to deal with, consider the movie database shown in Figure 1.² The first two subtrees out of the

root are both movies, but they differ in how they represent actors. The third subtree represents a TV series and has very little structure in common with movies. Furthermore, movies/TV series refer to each other, creating cycles in the structure. The situation is quite similar to that of a bibliographical database (cf. the examples given in [PGMW95]), where information about heterogeneous documents and their referencing relationship is represented in a graph-like structure.

Note that information only resides at labels. To express a structure in which information also resides at nodes (both ACeDB and Tsimmis allow this), we simply “migrate” that information down to a new edge attached to that node. Also note that there are three kinds of labels in the tree: character strings, integers and symbols—the latter corresponding to attribute names. Atomic data values such as strings and integers may occur anywhere in the tree, not just at terminal edges. This is similar to ACeDB, which allows integer labeled fields.

Outline of the paper. In the following sections we first develop the labeled tree data model and discuss correspondences with relational and nested relational databases. We then develop the query language UnQL (for Unstructured Query Language) through examples. On fixed-depth structures UnQL has the expressive power of the nested relational algebra. However, it can also express traversals of a tree of arbitrary depth, and these searches can be controlled by the use of regular expressions on paths. So far, this language is still limited in that it cannot *restructure* the database at arbitrary depths. We then give a more general—and more complex—construct that can restructure the database at arbitrary depths. After summarizing the

²With thanks to *The Internet Movie Database* (to be found at <http://www.msstate.edu/Movies/>). The deliberately irregular example used here in no way reflects the details of that well-structured and very useful database.

UnQL constructs, we give a calculus into which our languages can be translated and provide optimization techniques in this calculus. In particular we show that “horizontal” optimizations for relational structures can be used and that these have “vertical” counterparts that apply to deep queries. We conclude with some questions about the connection between UnQL and other languages.

2 The Data Model

Restricting ourselves for the moment to structures without cycles, we shall start with an explicit syntax for the construction of edge-labeled trees:

- $\{\}$ — the empty tree,
- $\{l \Rightarrow t\}$ — the tree whose root has an outgoing edge labeled l attached to the subtree t ,
- $t_1 \cup t_2$ — the union of trees t_1 and t_2 formed by coalescing the roots of t_1 and t_2 .

We shall also use the syntactic sugar $\{l_1 \Rightarrow t_1, l_2 \Rightarrow t_2, \dots, l_n \Rightarrow t_n\}$ for $\{l_1 \Rightarrow t_1\} \cup \{l_2 \Rightarrow t_2\} \cup \dots \cup \{l_n \Rightarrow t_n\}$.

For example, the tree

is specified by $\{A \Rightarrow \{B \Rightarrow \{\}, C \Rightarrow \{\}\}, D \Rightarrow \{\}\}$. We can further simplify the syntax by shortening the notation for terminal edges, $l \Rightarrow \{\}$, to l , and by omitting the braces around singleton trees $\{l \Rightarrow t\}$. Our example then becomes $\{A \Rightarrow \{B, C\}, D\}$.

Since labels carry all the basic information in this model, we need to specify the types of data that can be used as labels. We assume that the usual base types *String*, *Integer*, *Real*, etc., are available. These are the types that would be used as values in a relation. In addition we shall use a new type *Symbol* for labels that would correspond to attribute names in a relation. We write numbers and symbols literally (the latter usually capitalized) and use quotation marks for strings, e.g., “cat”. In what follows we make the simplifying assumption that labels can be symbols, strings, integers, etc; in fact, the type of labels is just the discriminated union of these base types.

Relational and nested relational databases.

Relational databases are easily encoded as trees. Starting with tuples, if a relation has attributes $A_1, A_2, \dots, A_n$, a tuple $\langle A_1 : v_1, A_2 : v_2, \dots, A_n : v_n \rangle$ may be encoded as $\{A_1 \Rightarrow v_1, A_2 \Rightarrow v_2, \dots, A_n \Rightarrow v_n\}$. Next, choosing a special label “*Tup*” to indicate the encoding of a tuple, we can encode a relation as $\{Tup \Rightarrow t_1, Tup \Rightarrow t_2, \dots, Tup \Rightarrow t_n\}$, where $t_1, t_2, \dots, t_n$ are the encodings of the tuples in the relation. Finally, the database itself is a set of relations named $R_1, R_2, \dots, R_n$, and the encoding of the database is $\{R_1 \Rightarrow r_1, R_2 \Rightarrow r_2, \dots, R_n \Rightarrow$

$r_n\}$, where $r_1, r_2, \dots, r_n$ are the encodings of the associated relations. For example, Figure 2 shows a simple relational database, first presented conventionally, then as a tree, and then in this encoding. Nested relational databases can be encoded similarly.

Cyclic structures. These can be built by extending our syntax for trees so that we can tie the leaf nodes back into the tree. We introduce *tree markers* $X_1, X_2, \dots$, and use a further syntactic construction:

$$e \text{ where } X_1 = e_1, X_2 = e_2, \dots, X_n = e_n.$$

This expression denotes a (possibly cyclic) “tree”. In it $e, e_1, e_2, \dots, e_n$ are tree expressions built up from the previous tree constructors augmented with tree markers, which may be used wherever a tree is used. For example,

$$\begin{aligned} X_1 \text{ where } X_1 &= \{A \Rightarrow X_2, C \Rightarrow \{D, E \Rightarrow X_1\}\}, \\ X_2 &= \{B \Rightarrow X_2\} \end{aligned}$$

denotes the cyclic structure depicted in Figure 3 (a). We shall continue to use the term “tree” to describe such structures, understanding that they are really rooted directed connected graphs.

Cyclic structures are easiest to draw with the aid of ϵ -edges, which are edges with no label. An ϵ -edge from node v to node w denotes that a copy of every edge emanating from w should be attached to v . For example, consider the following, rather tricky cyclic structure³: $X_1 \text{ where } X_1 = \{A \Rightarrow (X_1 \cup X_2)\}, X_2 = \{B \Rightarrow (\{C\} \cup X_1)\}$. We draw it, rather naturally, as in Figure 3 (b). Furthermore, we can prove that any graph with ϵ -edges is “equal” in a sense which is made precise through the notion of bisimulation [Mil89]) to one without ϵ -edges: the graph in Figure 3 (b) is equivalent to that in Figure 3 (c). See [Kos95, BDHS96] for details and related work.

Types. The type of cycle-free labeled trees has a simple description. Let *Label* be the type of edge labels and, for any type τ , let $\mathcal{P}_{\text{fin}}(\tau)$ describe the type of finite sets of τ . The type for labeled trees *Tree* satisfies the equation $\text{Tree} = \mathcal{P}_{\text{fin}}(\text{Label} \times \text{Tree})$.

Informally, this equation says that a tree is a set of pairs of labels and trees. Following [BTS91], we can obtain a natural form of computation for this type based on structural recursion. An initial version of this idea was pursued in [BDS94]. This paper describes new languages that are more general restrictions of structural recursion and are well-defined on cyclic structures; it also explores optimization techniques. In the development of these languages, it is important to remember that we are dealing with just two data types: *Label* and *Tree*.

³Recall that markers may be used wherever trees may be used: hence, $X_1 \cup X_2$ is a valid tree expression.

Figure 2: Representations of a relational database.

3 UnQL: Selection Queries

We begin by presenting some examples of UnQL that query the tree down to some fixed depth. On data structures that represent relational databases, such queries are equivalent to relational algebra queries. We then give examples that query to an arbitrary depth, and retrieve information from anywhere in the tree. Even though such queries contain complex “path expressions,” they are nevertheless meaningful for cyclic structures.

3.1 Scratching the Surface—the Top-Level Fragment of UnQL

Consider some simple examples on the relational database of Figure 2, which we assume is named *DB*.

Example 3.1 The expression

```
select t
where  $R1 \Rightarrow \backslash t \leftarrow DB$ 
```

says “compute the union of all trees *t* such that *DB* contains an edge $R1 \Rightarrow t$ emanating from the root.” There is only one such edge in Figure 2; this query therefore simply returns the set of tuples in *R1*. The returned expression is:

$$\{ Tup \Rightarrow \{A \Rightarrow "a", B \Rightarrow 2, C \Rightarrow 3\},$$

$$Tup \Rightarrow \{A \Rightarrow "b", B \Rightarrow 4, C \Rightarrow 5\}\}$$

There are some important differences between our *comprehension syntax* [BLS⁺94] and that of relational calculus or languages derived from SQL. In the *where* part of the expression (there is no *from* component), the form $R1 \Rightarrow \backslash t \leftarrow DB$ is a *generator*, which has

two components, a *pattern* $R1 \Rightarrow \backslash t$ (everything to the left of the “generator” ($\leftarrow$) arrow) and a tree *DB* (the right-hand side of the $\leftarrow$ arrow). A tree is simply a set of edge/subtree pairs, so the generator matches the pattern to each of these pairs. Only those edges labeled *R1* will match. When such a match occurs, the variable $\backslash t$ is bound to the associated subtree, and the expression in the *select* clause is evaluated and added to the result.

When variables are introduced in UnQL, they are flagged with a backslash. They can then be used later (without the backslash) in the *where* part of the comprehension as well as in the *select* part. This explicit binding of variables is needed to avoid ambiguities that would arise in nested queries.

Example 3.2 The following expression uses a label variable $\backslash l$ to match any edge emanating from the root.

```
select t
where  $\backslash l \Rightarrow \backslash t \leftarrow DB$ 
```

The result is the union of all tuples in both relations—a heterogeneous set that cannot be described by a single relation.

Example 3.3 Here we join *R1* and *R2* on their common attribute *C* and then project onto *A* and *D*:

```
select  $\{Tup \Rightarrow \{A \Rightarrow x, D \Rightarrow z\}\}$ 
where  $R1 \Rightarrow Tup \Rightarrow \{A \Rightarrow \backslash x, C \Rightarrow \backslash y\} \leftarrow DB,$ 
 $R2 \Rightarrow Tup \Rightarrow \{C \Rightarrow y, D \Rightarrow \backslash z\} \leftarrow DB$ 
```

$R1 \Rightarrow Tup \Rightarrow \{A \Rightarrow \backslash x, C \Rightarrow \backslash y\}$ is a *tree pattern*. It matches any subtree of *DB* starting at the root, continuing with an *R1* edge, then with a *Tup* edge which

Figure 3: Cyclic structures.

has two edges labeled A and C respectively. Moreover it binds x and y to the subtrees after A and C . Note that the variable y is bound in the pattern of one generator and then used as a constant in the pattern of the second.

Example 3.4 This query performs a group-by operation on $R2$ along the C column:

```
select {x ⇒ (select y where
    R2 ⇒ Tup ⇒ {C ⇒ x, D ⇒ \y} ← DB)}
where R2 ⇒ Tup ⇒ C ⇒ \x ⇒ {} ← DB
```

The use of $\backslash x \Rightarrow \{\}$ is needed to bind x to an edge label rather than a tree, so that the expression $x \Rightarrow \dots$ in the `select` clause makes sense. In contrast, $\backslash y$ ranges over trees. The output from this last query is the tree $\{3 \Rightarrow \{"c"\}, 5 \Rightarrow \{"d", "e"\}\}$.

Example 3.5 Turning to the movie database in Figure 1, we see that there are some queries that can be answered by the techniques we have already developed. For example, give the titles and casts of all entries—movies, TV shows, etc.

```
select {Tup ⇒ {Title ⇒ x, Cast ⇒ y}}
where Entry ⇒ _ ⇒ {Title ⇒ \x, Cast ⇒ \y} ← DB
```

Here the “wildcard” symbol $_$ matches any edge label; its use is equivalent to binding a fresh anonymous variable at that point. Also note that the query has returned a tree that is a set of pairs, but the components of these pairs (x and y) are trees.

Example 3.6 The next example is more problematic. We want a binary relation consisting of actress/actor and title tuples for movies. The problem is that the information is not uniformly located within the tree.

```
select {Tup ⇒ {Actor ⇒ x, Title ⇒ y}}
where
    Entry ⇒ Movie ⇒ {Title ⇒ \y, Cast ⇒ \z} ← DB,
    \x ⇒ {} ← z ∪ (select u where _ ⇒ \u ← z),
    isstring(x)
```

The assumption we have made is that the names we want will be found immediately below the *Cast* edge or one step further down. Moreover we only want those

edges that are strings. This query illustrates the use of a *condition*, which is the second construct that can occur in the `where` part of a comprehension. We also assume that the type of labels is a discriminated union, so properties such as *isstring* should be available.

3.2 Taking the Plunge: Deep Queries.

It is apparent from the previous examples that we need more expressive power if we are to look for data whose depth in the tree is not predetermined by any schema.

Example 3.7 Perhaps the simplest of all queries that look arbitrarily deep into the database is to find all edges with a certain property. For example, to find the set of all strings in the database:

```
select {l}
where _* ⇒ \l ⇒ _ ← DB, isstring(l)
```

Here the $_*$ is a “repeated wildcard” that matches any path, i.e., sequence of edges, in the tree. Such a construct is proposed in [PGMW95]. However, we shall find some useful queries where we need more than this, i.e. we need to specify regular expressions on paths. For this, we adopt a *grep*-like syntax.

One may wonder whether queries containing a $_*$ construct are well defined on cyclic structures, given that the number of paths in such structures is infinite. However, since cyclic structures have only a finite number of distinct subtrees, there is only a finite number of distinct assignments of labels and trees to the variables in the `where` clause, so the output of the query is still finite.

The use of a leading $_*$ is so common that we shall use a special abbreviation $p \leftarrow t$ for $_* \Rightarrow p \leftarrow t$ for any pattern p and any tree expression t . The query above then becomes:

```
select {l}
where \l ⇒ _ ← DB, isstring(l)
```

Example 3.8 Here we use consecutive “deep” generators to find all the movies involving “Bogart” and “Bacall”:

```
select {Movie ⇒ x}
where Movie ⇒ \x ← DB,
```

"Bogart" $\Rightarrow$ $_ \leftarrow x$,

"Bacall" $\Rightarrow$ $_ \leftarrow x$

This will find all movie edges in the database (no matter where they are) and return those edges, with their subtrees, that contain both "Bogart" and "Bacall" somewhere beneath them.

Example 3.9 A problem with the example above is that it may return more than is required. For example, a movie will be returned if it *refers* to a movie in which Bogart and Bacall are involved. In order to avoid such paths we use a more complicated pattern in a path:

```
select {Movie  $\Rightarrow$  x}
where Movie  $\Rightarrow$  \x  $\leftarrow$  DB,
      [^Movie]*  $\Rightarrow$  "Bogart"  $\Rightarrow$  \_  $\leftarrow$  x,
      [^Movie]*  $\Rightarrow$  "Bacall"  $\Rightarrow$  \_  $\leftarrow$  x
```

Following `grep`, the pattern $[^Movie]^*$ matches any path that does not contain the label *Movie*. Arbitrary regular expressions may be used on labels.

4 UnQL: Restructuring Queries

All the queries discussed so far have only exploited depth by pulling out subtrees. In this section we describe the *traverse* query construct, whose output tree can represent arbitrarily deep changes in the input tree. Consider first a simple example:

Example 4.1 Replace all *SpecialGuest* labels with a *Featuring* label:

```
traverse DB giving X
case SpecialGuest  $\Rightarrow$  \_ then X := {Featuring  $\Rightarrow$  X}
case \l  $\Rightarrow$  \_          then X := {l  $\Rightarrow$  X}
```

The construction says that we have to replace every edge of the tree *DB* with some tree as follows. Whenever we see an edge labeled *SpecialGuest* in *DB*, we replace it with the tree $\{Featuring \Rightarrow X\}$: in essence this is just an edge, leading to a marker *X*. Whenever we see an edge labeled something else, like *l*, we replace it with the tree $\{l \Rightarrow X\}$. The *X*'s of the different trees are now joined as follows: each *X*-leaf of some tree replacing an edge *e* is joined with the roots of the trees replacing the successors of *e*. The order in which we write the two *case* statements matters: had we reversed them, we would have produced a tree identical to the original one.

The restructuring performed by the *traverse* construct is relatively simple, in that it acts only locally—on edges. This simplicity allows it to have meaning on cyclic structures (cycles will simply be transformed into other cycles), but limits its expressive power. We can increase its expressiveness significantly by allowing several markers instead of one, while still confining ourselves to local restructurings, as illustrated in the following example.

Figure 4: Replacing all *SpecialGuest* edges underneath *TVShows*. The new root is the root of the first tree. Note that part of both trees are inaccessible from this root: they need not be constructed during query evaluation.

Example 4.2 Suppose we want to replace every edge *SpecialGuest* with a *Featuring* edge, but only in *TVShows*, and leave the *SpecialGuest* edge for other kinds of movies untouched. Being part of a *TVShow* is not local information available at that edge, and in general we don't know how far apart the *TVShow* edge is from its underlying *SpecialGuest*. The trick here is to construct two trees in parallel, instead of one, having edges going from the first to the second one, as illustrated in Figure 4. The first tree, t_1 , will be an identical copy of the input database *DB*. The second tree, t_2 , will be a copy of *DB* with all *SpecialGuest* edges replaced by *Featuring*. In addition, every edge labeled *TVShow* in t_1 will be redirected to the corresponding node in t_2 . As a consequence, every edge in *DB* will be replaced by two "small" trees, and we will use two markers X_1 and X_2 to keep track of them:

```
traverse DB giving X1, X2
case TVShow  $\Rightarrow$  \_ then X1 := {TVShow  $\Rightarrow$  X2}
case SpecialGuest  $\Rightarrow$  \_ then X2 := {Featuring  $\Rightarrow$  X2}
case \l  $\Rightarrow$  \_ then X1 := {l  $\Rightarrow$  X1},
                  X2 := {l  $\Rightarrow$  X2}
```

In this construction, pattern matching is done twice for each edge, once for X_1 and once for X_2 . Clauses that do not provide an assignment to the marker under consideration are disregarded. Thus, on a *TVShow* edge, the first clause would succeed for X_1 and the last clause for X_2 .

By convention, the final result is the tree rooted at the first marker X_1 , that is, t_1 . Parts of t_1 and t_2 will be superfluous, like the edges in t_1 beneath a *TVShow* edge, and those in t_2 above a *TVShow* edge. Since they are inaccessible from the root they can be discarded; in

fact, during the execution of this query those parts of t_1 and t_2 need not even be constructed.

Example 4.3 Consider the following deep-nest example: in a given tree we want to nest all edges labeled A . That is, every subtree of the form $\{A \Rightarrow t_1, A \Rightarrow t_2, \dots, A \Rightarrow t_n\} \cup t$, where t has no top-most A edges, has to be replaced with $\{A \Rightarrow t_1 \cup \dots \cup t_n\} \cup t$. The query is given below: for simplicity we assume that no A -edges occur at the topmost level and that there are no two consecutive A -edges in the tree.

```

traverse DB giving  $X_1, X_2$ 
  case  $A \Rightarrow \_$       then  $X_1 := \{\}, X_2 := X_1$ 
  case  $\backslash l \Rightarrow A \Rightarrow \_$  then  $X_1 := \{l \Rightarrow (X_1 \cup \{A \Rightarrow X_2\})\},$ 
                                 $X_2 := \{\}$ 
  case  $\backslash l \Rightarrow \_$       then  $X_1 := \{l \Rightarrow X_1\}, X_2 := \{\}$ 

```

traverse and select are typically executed by traversing the tree and, in the process, constructing the resulting tree. When two such queries are composed, it might appear that one has to traverse the whole tree, build an intermediate result, and then traverse the result. However, we show in Section 6 that the composed query can be optimized: we “fuse” together the two traversals into a single one and eliminate the intermediate result. For example, recall the query in Example 4.2 returning a tree T with some of its *SpecialGuest* edges replaced by *Featuring* edges. Now assume that we want to compose this query with: `select {SpecialGuest  $\Rightarrow$  t} where SpecialGuest  $\Rightarrow$  t  $\leftarrow$  T`. By composing the two queries we obtain the set of all *SpecialGuests*, but excluding those from *TVShows*. Of course this can be easily expressed as a single `select` query. We will show in Section 6 how this single query can be derived from simple algebraic equations.

5 Summary of UnQL

For completeness, we briefly enumerate UnQL’s operators: `{}`, `{ $\_ \Rightarrow \_$ }`, `$\cup$` , `@`, `if  $\_$  then  $\_$  else  $\_$` , `select`, `traverse`. We have already seen examples of all these with the exception of `@`, which is described in Section 6; see [BDHS96] for a more complete description of `traverse`. Here we summarize the syntax for `select`:

```

select E
where  $C_1, C_2, \dots, C_m$ 

```

E is another UnQL expression and each C_i is either a *generator* or a *condition*. Conditions are predicates on labels such as $l_1 = l_2$ or $isstring(l)$, and tree emptiness tests like $isempty(t)$. Generators are of the form $P \leftarrow T$, where P is a pattern and T is a UnQL expression. A pattern P is either (1) a tree constant or a tree variable, or (2) $\{R_1 \Rightarrow P_1, \dots, R_m \Rightarrow P_m\}$ with each P_i a pattern, and each R_i either a variable $\backslash l$, or a regular expression. We abbreviate $\{R_1 \Rightarrow P_1\}$ with $R_1 \Rightarrow P_1$.

Regular expressions may use, but not bind variables, hence we can have $\backslash l \Rightarrow \backslash t \leftarrow DB$ and $l* \Rightarrow \backslash t \leftarrow DB$, but not $(\backslash l)* \Rightarrow \backslash t \leftarrow DB$. Some examples of more complex patterns are: $(_ \Rightarrow _)* \Rightarrow \backslash l \Rightarrow \backslash t \leftarrow DB$ binds l to an edge label which is at an even depth in DB . $[\textit{TVShow}]* \Rightarrow \backslash l \Rightarrow [\textit{SpecialGuest}]* \Rightarrow l \Rightarrow \backslash t \leftarrow DB$ binds l to an edge label which occurs at least twice in DB , first before any *TVShow* edge and second before any *SpecialGuest* edge.

Recall that the tree model naturally encodes relational databases as trees of depth 4 (Figure 2). Similarly, it encodes nested relational databases as trees of a fixed depth. On these inputs, UnQL has the same expressive power as relational algebra, or nested relational algebra.

Theorem 5.1 *On trees encoding relational databases, or nested relational databases, UnQL has exactly the same expressive power as relational algebra, or nested relational algebra respectively.*

However, note that UnQL can be more succinct than relational algebra: the UnQL query

```

select {Tup  $\Rightarrow$  t}
where  $\_ \Rightarrow$  Tup  $\Rightarrow$  t  $\leftarrow$  DB,
       $\_ \Rightarrow$  "Smith"  $\leftarrow$  t

```

selects all tuples containing the string “Smith” from the database. Ignoring for the moment the fact that the result is a heterogeneous set, one can translate this query into relational algebra, but there will be a *different* translation for each database schema: as the schema becomes more complex, so does the translated query.

6 A Calculus for UnQL and its Optimizations

We now show that UnQL is equivalent to a simple calculus, called UnCAL for *unstructured calculus*. Following [BBW92], we use the term *calculus* in the sense of *lambda calculus* (a formalism with variables and functions), rather than in the sense of *relational calculus* (a logic, with variables and quantifiers). We sketch a simple, effective procedure for translating UnQL queries into UnCAL expressions. By simplifying UnQL into UnCAL, we achieve two purposes. First, we can prove properties about UnQL, for example that all queries are defined even on cyclic structures and are computable in PTIME. Second, and more important, we derive optimization techniques for UnQL based on the algebraic identities that hold for UnCAL. In this section we show some of these equations and illustrate the corresponding optimizations.

6.1 The Calculus UnCAL

The salient construct in UnCAL is its recursion operator *gext*, which is the mechanism underlying the `traverse`

construct. Like *traverse*, this is a rather unorthodox operator, in that it uses tree markers in an essential way. While real-life examples probably will not put too much burden on these markers, the interactions between different sets of markers that result during the optimization steps are, as we shall see shortly, somewhat delicate. Hence UnCAL is quite picky about keeping track of the markers. For a finite set $\mathcal{X}$ of markers, we will denote by $Tree_{\mathcal{X}}$ the set of trees having only markers in $\mathcal{X}$. Then, $\mathcal{X} \subseteq \mathcal{X}'$ implies $Tree_{\mathcal{X}} \subseteq Tree_{\mathcal{X}'}$. Furthermore, for a reason that will become apparent shortly, we will assume that we can put together two markers to build a new one; that is, $X \cdot Y$ will be a new marker built from X and Y , and it will be distinct from all other markers built in this way. For two sets of markers $\mathcal{X}$ and $\mathcal{Y}$, $\mathcal{X} \cdot \mathcal{Y}$ denotes the set of markers $X \cdot Y$, with $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$.

UnCAL's operators, which are listed in Figure 5, exploit the dual nature of our trees: *sets* along their horizontal dimension and *branching lists* along the vertical one. UnCAL has set-like operators for the horizontal dimension of trees and list-like operators for the vertical dimension. This is best illustrated by the union operator $\cup$, which joins two trees horizontally, and the append operator $@$, which joins them vertically. As expected, append is more complex than union. In an append expression, $t @_{\mathcal{X}} \vec{s}$, t is a tree of type $Tree_{\mathcal{X}}$ and $\vec{s}$ is an $\mathcal{X}$ -indexed family of trees of type $Tree_{\mathcal{Y}}$. Such families occur often in UnCAL, and we write them as $\vec{s} = \{X_1 := s_1, \dots, X_n := s_n\}$, when $\mathcal{X} = \{X_1, \dots, X_n\}$; we denote the set of $\mathcal{X}$ -indexed families of trees with $\mathcal{Y}$ markers with $Tree_{\mathcal{Y}}^{\mathcal{X}}$. The result of append is obtained by replacing each marker X_i in t with the tree s_i . The $\mathcal{Y}$ markers in the s_i are preserved in the output, which is therefore of type $Tree_{\mathcal{Y}}$.

As a simple example with only one marker, consider $t = \{A \Rightarrow B, C \Rightarrow X\} \cup X$ and⁴ $s = \{B, D \Rightarrow E\}$. Then $t @_{\{X\}} s = \{A \Rightarrow B, C \Rightarrow \{B, D \Rightarrow E\}, B, D \Rightarrow E\}$. For an example with two markers, let $t = \{A \Rightarrow (X_1 \cup X_2), B \Rightarrow (\{C\} \cup X_1)\}$, and let $\vec{s}$ be the $\{X_1, X_2\}$ -indexed family $\{X_1 := \{B, D\}, X_2 := \{B, D \Rightarrow E\}\}$. Then $t @_{\{X_1, X_2\}} \vec{s} = \{A \Rightarrow \{B, D, D \Rightarrow E\}, B \Rightarrow \{C, B, D\}\}$. See Figure 6 (a) and (b).

UnCAL's most complex operator is $gext_{\mathcal{X}}(f)(T)$. Here f is a function taking a label l and a tree t as inputs and returning an $\mathcal{X}$ -indexed family of trees with markers $\mathcal{X}$, while the result is an $\mathcal{X}$ -indexed family of trees. $gext_{\mathcal{X}}(f)(T)$ has the same meaning as a certain restricted form of *traverse*:

traverse T giving $\mathcal{X}$
 case $\backslash l \Rightarrow \backslash t$ then $f(l, t)$

We will describe the semantics of *gext* in detail next. First we consider the case where $\mathcal{X}$ has only one marker,

⁴Strictly speaking s is the $\{X\}$ -indexed family $\{X := \{B, D \Rightarrow E\}\}$; we omit mentioning X when it is clear from the context.

$$\begin{array}{c}
 \frac{}{\{\} : Tree_{\mathcal{X}}} \quad \frac{X \in \mathcal{X}}{X : Tree_{\mathcal{X}}} \quad \frac{l : Label \quad t : Tree_{\mathcal{X}}}{\{l \Rightarrow t\} : Tree_{\mathcal{X}}} \\
 \\
 \frac{l_1 : Label \quad l_2 : Label}{l_1 = l_2 : Bool} \quad \frac{t : Tree_{\mathcal{X}}}{isempty(t) : Bool} \\
 \\
 \frac{b : Bool \quad t_1 : Tree_{\mathcal{X}} \quad t_2 : Tree_{\mathcal{X}}}{if \ b \ then \ t_1 \ else \ t_2 : Tree_{\mathcal{X}}} \\
 \\
 \frac{t_1 : Tree_{\mathcal{X}} \quad t_2 : Tree_{\mathcal{X}}}{t_1 \cup t_2 : Tree_{\mathcal{X}}} \quad \frac{t : Tree_{\mathcal{X}} \quad \vec{s} : Tree_{\mathcal{Y}}^{\mathcal{X}}}{t @_{\mathcal{X}} \vec{s} : Tree_{\mathcal{Y}}} \\
 \\
 \frac{f : Label \times Tree_{\mathcal{Y}} \rightarrow Tree_{\mathcal{X}}^{\mathcal{X}} \quad t : Tree_{\mathcal{Y}}}{gext_{\mathcal{X}}(f)(t) : Tree_{\mathcal{X}}^{\mathcal{X}}}
 \end{array}$$

Figure 5: The rules for UnCAL.

Figure 6: Two examples of append, $t @ t'$.

Figure 7: Visualization of $\text{gext}(f)(T)$: f is applied independently to each edge, and the resulting “small” trees are joined.

say X . Then we describe gext as a *recursive process*. Let $g(T) = \text{gext}_{\{X\}}(f)(T)$, with T a tree having at most markers $Y_1, \dots, Y_n$. Then:

$$\begin{aligned} g(\{\}) &\stackrel{\text{def}}{=} \{\} \\ g(\{l \Rightarrow t\}) &\stackrel{\text{def}}{=} f(l, t) @_{\{X\}} g(t) \\ g(T_1 \cup T_2) &\stackrel{\text{def}}{=} g(T_1) \cup g(T_2) \\ g(Y_i) &\stackrel{\text{def}}{=} Y_i \text{ for } i = 1, n \end{aligned}$$

This suggests a computation method in which we start from the root and go downwards in the tree. But one wonders if gext loops forever on cycles. The answer is no: one can show [BDHS96] that on a cyclic structure $T = (Y_1 \text{ where } Y_1 = T_1, \dots, Y_n = T_n)$, $g(T)$ is $Y_1 \text{ where } Y_1 = g(T_1), \dots, Y_n = g(T_n)$.

Alternatively, it is useful to view gext as a *parallel process*. Namely to compute $\text{gext}_{\{X\}}(f)(T)$, we start by making a copy v' for every node v in the tree T . Next, we take each edge $l \Rightarrow t$ in T going from vertex v to w , and replace it with $f(l, t)$. When doing this, we identify⁵ the root of $f(l, t)$ with the vertex v' and identify all leaves labeled X with w' . See Figure 7 for an illustration.

We can prove [BDHS96] that the recursive and parallel processes describing gext are equivalent. The parallel one makes it clear that the computation does not go into an infinite loop on a cyclic structure and justifies the somewhat arbitrary equation $g(Y_i) \stackrel{\text{def}}{=} Y_i$. However, the recursive view is sometimes more appropriate when designing queries.

Now we turn to the general case, when $\mathcal{X} = \{X_1, \dots, X_m\}$. Then $\text{gext}_{\mathcal{X}}$ allows us to define simultaneously m mutually recursive functions g_i , $i = 1, m$, as shown in Figure 8. When we apply $\text{gext}_{\mathcal{X}}(f)$ to a tree with n markers $Y_1, \dots, Y_n$, we get a tree with mn markers, labeled $X_i \cdot Y_j$, $i = 1, m$, $j = 1, n$.

We specialize $\text{gext}_{\mathcal{X}}(f)$ to two important cases. The

first is when $\mathcal{X} = \emptyset$.⁶ In this case only the topmost level of the tree denoted by t is processed. This is a classical operation on sets which we call ext :

$$\frac{f : \text{Label} \times \text{Tree} \rightarrow \text{Tree} \quad t : \text{Tree}}{\text{ext}(f)(t) : \text{Tree}}$$

with the meaning $\text{ext}(f)(\{l_1 \Rightarrow t_1, \dots, l_n \Rightarrow t_n\}) = f(l_1, t_1) \cup \dots \cup f(l_n, t_n)$. The second case is when $f(l, t)$ depends only on the label l , not on the tree t . We call this vext .

$$\frac{f : \text{Label} \rightarrow \text{Tree}_{\mathcal{X}} \quad t : \text{Tree}_{\mathcal{Y}}}{\text{vext}_{\mathcal{X}}(f)(t) : \text{Tree}_{\mathcal{X} \cdot \mathcal{Y}}}$$

When $\mathcal{X}$ has a single marker X , then vext is exactly ext on lists. Indeed, assume that we encode a list $[A_1, \dots, A_n]$ “vertically” as $\{A_1 \Rightarrow \{A_2 \Rightarrow \dots \{A_n\} \dots\}\}$, and assume that $f(l)$ returns a list “ending in X ” (i.e., of the form $\{B_1 \Rightarrow \{B_2 \Rightarrow \dots \{B_m \Rightarrow X\} \dots\}\}$), for every label l . Then vext is described by $\text{vext}(f)([A_1, \dots, A_n]) \stackrel{\text{def}}{=} f(A_1) @ \dots @ f(A_n)$.

An example of vext is the query $\text{deepflatten}(T)$, which returns the set of all subtrees of t . In UnQL, this is expressed as

```
select t
where _* => \t ← T
```

In UnCAL we express it as the X_1 -component of:

$$\text{vext}_{\{X_1, X_2\}}(\lambda l. X_1 := X_1 \cup \{l \Rightarrow X_2\}, X_2 := \{l \Rightarrow X_2\})(T)$$

More generally, for any deterministic finite automaton there exists a $\text{vext}_{\mathcal{X}}$ expression that pulls up to the top level all subtrees reachable by a path accepted by that automaton. Moreover, the automaton may have not only labels, but also previously bound variables, or even negation (i.e., $\neq l$) on its edges. The number of markers in $\mathcal{X}$ is one plus the number of states of the automaton. In particular, all queries $\text{select } t \text{ where } R \Rightarrow \backslash t \leftarrow T$ from Section 3, where R is a regular expression, can be translated into vext .

Theorem 6.1 *UnQL and UnCAL express the same set of queries. Moreover, there exists an effective translation from UnQL to UnCAL.*

UnCAL is a simpler formalism than UnQL. Hence we can more easily prove properties about UnCAL than about UnQL. For example:

Theorem 6.2 *All queries expressible in UnCAL (and, hence, in UnQL) are defined on all cyclic structures and are computable in PTIME.*

⁵Equivalently, join them with ϵ edges.

⁶By convention, in this case f returns a single tree, and $\text{gext}_{\mathcal{X}}(f)$ returns a single tree too.

$$\begin{array}{llll}
g_1(\{\}) & \stackrel{\text{def}}{=} & \{\} & \dots & g_m(\{\}) & \stackrel{\text{def}}{=} & \{\} \\
g_1(\{l \Rightarrow t\}) & \stackrel{\text{def}}{=} & f_1(l, t) @_{\mathcal{X}} g(t) & \dots & g_m(\{l \Rightarrow t\}) & \stackrel{\text{def}}{=} & f_m(l, t) @_{\mathcal{X}} g(t) \\
g_1(T_1 \cup T_2) & \stackrel{\text{def}}{=} & g_1(T_1) \cup g_1(T_2) & \dots & g_m(T_1 \cup T_2) & \stackrel{\text{def}}{=} & g_m(T_1) \cup g_m(T_2) \\
g_1(Y_j) & \stackrel{\text{def}}{=} & X_1 \cdot Y_j & \dots & g_m(Y_j) & \stackrel{\text{def}}{=} & X_m \cdot Y_j
\end{array} \tag{1}$$

Figure 8: The definition of $\text{gext}_{\mathcal{X}}(f)(T)$. Here $f(l, t)$ and $g(t)$ are $\mathcal{X}$ -indexed families $\{X_1 := f_1(l, t), \dots, X_m := f_m(l, t)\}$, $\{X_1 := g_1(t), \dots, X_m := g_m(t)\}$ and $g(T) = \text{gext}_{\mathcal{X}}(f)(T)$.

6.2 Optimizations in UnCAL

The main advantage of having a simple formalism is that it enables us to derive optimization rules. Optimization rules for set-based query languages centered around ext and set union $\cup$ have been considered previously [Won94]. Surprisingly, they hold in identical form for the vertical dimension, with the slight twist that sometimes we have to restrict ourselves to vext instead of gext . Rather than considering all optimization rules, we shall focus on two: the first is a simple equation pushing selections down a join tree, the second a rather complex loop fusion, putting maximum strain on the usage of markers in UnCAL.

Example 6.3 We start with the following simple optimization rule for ext , adapted from [Won94]:

$$\begin{aligned}
& \text{ext}(\lambda(l, t). \text{if } p \text{ then } e \text{ else } e')(T) = \\
& \quad \text{if } p \text{ then } \text{ext}(\lambda(l, t).e)(T) \text{ else } \text{ext}(\lambda(l, t).e')(T) \\
& \quad (\text{when } l \text{ and } t \text{ do not occur free in } p)
\end{aligned}$$

This extends naturally to vext and gext as:

$$\begin{aligned}
& \text{gext}(\lambda(l, t). \text{if } p \text{ then } e \text{ else } e')(T) = \\
& \quad \text{if } p \text{ then } \text{gext}(\lambda(l, t).e)(T) \text{ else } \text{gext}(\lambda(l, t).e')(T) \\
& \quad (\text{when } l \text{ and } t \text{ do not occur free in } p)
\end{aligned}$$

Using this, we derive the classical optimization rule pushing selections down a join tree for the vertical dimension of UnCAL. To see how this works, consider the following *vertical selection-join* query:

$$\begin{aligned}
& \text{traverse } t_1 \text{ giving } X \\
& \quad \text{case } \backslash a \Rightarrow _ \text{ then} \\
& \quad \quad X := (\text{traverse } t_2 \text{ giving } Y \\
& \quad \quad \quad \text{case } \backslash b \Rightarrow _ \text{ where } p(a) \text{ then} \\
& \quad \quad \quad \quad Y := \{f(a, b) \Rightarrow Y\})
\end{aligned}$$

Here each edge a of t_1 satisfying $p(a)$ is replaced with a full copy of the tree t_2 ; edges not satisfying $p(a)$ are removed. The copy of t_2 is obtained by replacing each of its labels b with a new label $f(a, b)$. To make the query nontrivial, we assume that t_2 had initially some markers X . The query gets translated into the following UnCAL expression: $\text{vext}_{\{X\}}(\lambda a. \text{vext}_{\{Y\}}(\lambda b. \text{if } p(a) \text{ then } \{f(a, b) \Rightarrow Y\} \text{ else } \{\}) (t_2)) (t_1)$.

The problem is that the predicate p is applied in the inner loop, while it would be obviously less expensive to apply it in the outer loop, as in:

$$\begin{aligned}
& \text{traverse } t_1 \text{ giving } X \\
& \quad \text{case } \backslash a \Rightarrow _ \text{ where } p(a) \text{ then} \\
& \quad \quad X := (\text{traverse } t_2 \text{ giving } Y \\
& \quad \quad \quad \text{case } \backslash b \Rightarrow _ \text{ then} \\
& \quad \quad \quad \quad Y := \{f(a, b) \Rightarrow Y\})
\end{aligned}$$

which is translated into $\text{vext}_{\{X\}}(\lambda a. \text{if } p(a) \text{ then } \text{vext}_{\{Y\}}(\lambda b. \{f(a, b) \Rightarrow Y\}) (t_2) \text{ else } \{\}) (t_1)$. The rule for if described earlier and the simple rule $\text{vext}(f)(\{\}) = \{\}$ allows us to transform the first query into the second one.

The second optimization rule corresponds to the *vertical loop fusion* in [GP84]. Consider the expression $\text{vext}(g)(\text{vext}(f)(t))$. One can visualize it as a pipeline process in which the edges of t are fed one by one into the function f , which transforms every label into a “small” tree $f(a)$. The new edges b produced by f are further fed into g , which replaces each such edge with another “small” tree $g(b)$. One can optimize this process by traversing t only once, and replacing every edge a with the tree $\text{vext}(g)(f(a))$. This makes the intermediate result $\text{vext}(f)(t)$ unnecessary, and may even speed up the computation dramatically (e.g., when some of the $g(b)$ trees are empty) by pruning early the computation of $\text{vext}(f)$. In short, we perform an optimization by replacing $\text{vext}(g)(\text{vext}(f)(t))$ with $\text{vext}(\text{vext}(g) \circ f)(t)$ (here $\circ$ denotes function composition).

Getting this rule to work well was difficult until we discovered the way in which $\text{vext}_{\mathcal{X}}(f)(t)$ must interact with markers existing in t , which is essentially given in equation (1). Before listing the optimization rules, let us consider an example.

Example 6.4 Take the query in Example 4.2, which replaces all *SpecialGuest* edges with *Featuring* edges in the *TVShows* of the movie database, and suppose this query defines a view T . T can be expressed as $\text{vext}_{\{X_1, X_2\}}(f)(DB)$, where f is graphically represented in Figure 9. Now assume that we want to query the view to obtain the set of all *SpecialGuests* in T . This

Figure 9: Graphic representation of the functions f and g .

Figure 10: Graphic representation of the function $h = \text{vext}(g) \circ f$.

can be written as $\text{select } \{ \text{SpecialGuest} \Rightarrow t \}$ where $_ \Rightarrow \text{SpecialGuest} \Rightarrow _ \mid t \in T$ or, equivalently, as:

```

traverse T giving Y1, Y2
case SpecialGuest  $\Rightarrow$  _ then
  Y1 := {SpecialGuest  $\Rightarrow$  Y2}
case \a  $\Rightarrow$  _ then
  Y1 := Y1, Y2 := {a  $\Rightarrow$  Y2}

```

This is precisely $\text{vext}_{\{Y_1, Y_2\}}(g)(T)$, where the function g is represented graphically in Figure 9. Querying the view T entails computing the composition $\text{vext}_{\{Y_1, Y_2\}}(g)(\text{vext}_{\{X_1, X_2\}}(f)(DB))$. This is a pipeline process as described earlier, and is equivalent to $\text{vext}_{\mathcal{Z}}(h)(DB)$, where $\mathcal{Z} = \mathcal{Y} \cdot \mathcal{X} = \{Y_1 \cdot X_1, Y_2 \cdot X_1, Y_1 \cdot X_2, Y_2 \cdot X_2\}$, and h is obtained by applying g locally on the “small” trees produced by f . The best way to understand the function h is graphically (see Figure 10). Nevertheless, we can write down the resulting query, consisting of a single **traverse** expression:

```

traverse T giving Y1 · X1, Y2 · X1, Y1 · X2, Y2 · X2
case TVShow  $\Rightarrow$  _ then
  Y1 · X1 := Y1 · X2,
  Y2 · X1 := {TVShow  $\Rightarrow$  Y2 · X2}
case SpecialGuest  $\Rightarrow$  _ then
  Y1 · X1 := {SpecialGuest  $\Rightarrow$  Y2 · X1},
  Y2 · X2 := {Featuring  $\Rightarrow$  Y2 · X2}
case \l  $\Rightarrow$  _ then
  Y1 · X1 := Y1 · X1,
  Y1 · X2 := Y1 · X2,
  Y2 · X1 := {\l  $\Rightarrow$  Y2 · X1},
  Y2 · X2 := {\l  $\Rightarrow$  Y2 · X2}

```

Further optimizations are possible by “reducing” the number of states (markers): in our example the tree corresponding to $Y_1 \cdot X_2$ is empty, because the only edge leaving from this “state” is an ϵ -edge into the same state, hence this marker can be eliminated.

We give next the optimization rules in their general form. As mentioned earlier, all 9 rules in [Won94] apply to the vertical operators in UnCAL. Instead of listing all of them, we show only the most powerful ones, together with their corresponding horizontal versions.

Theorem 6.5 *The following equations hold:*

$$\begin{aligned}
\text{ext}(f)(t_1 \cup t_2) &= \text{ext}(f)(t_1) \cup \text{ext}(f)(t_2) \\
\text{vext}_{\mathcal{X}}(f)(t_1 @_{\mathcal{Y}} t_2) &= \text{vext}_{\mathcal{X}}(f)(t_1) @_{\mathcal{X} \cdot \mathcal{Y}} \text{vext}_{\mathcal{X}}(f)(t_2) \\
\text{ext}(g)(\text{ext}(f)(t)) &= \text{ext}(\text{ext}(g) \circ f)(t) \\
\text{vext}_{\mathcal{X}}(g)(\text{gext}_{\mathcal{Y}}(f)(t)) &= \text{gext}_{\mathcal{X} \cdot \mathcal{Y}}(\text{vext}_{\mathcal{X}}(g) \circ f)(t)
\end{aligned}$$

The second equation emphasizes the parallel nature of vext . We can compute $\text{vext}_{\mathcal{X}}(f)$ on $t_1 @_{\mathcal{Y}} t_2$ by applying it independently on t_1 and t_2 (note that this fails for gext), and then appending the two results. Note

how the markers interact here: t_1 has markers $\mathcal{Y}$, while t_2 is a $\mathcal{Y}$ -indexed family of trees. Then $\text{vext}_{\mathcal{X}}(f)(t_1)$ has markers in $\mathcal{X} \cdot \mathcal{Y}$, hence the index in @ on the right hand side. Since t_2 is a $\mathcal{Y}$ -indexed family of trees, $\text{vext}_{\mathcal{X}}(f)(t_2)$ is a $\mathcal{X} \cdot \mathcal{Y}$ -indexed family of trees, obtained by aggregating all $\mathcal{X}$ -indexed families of trees obtained by applying $\text{vext}_{\mathcal{X}}(f)$ to each tree in the $\mathcal{Y}$ -family t_2 .

The last equation is the vertical loop fusion in its most general form: Example 6.4 is an instance of this equation. Note that the loop fusion works even when the first loop is a *gext*: the second loop however needs to be a *vext*. Also note that the new loop function, $\text{vext}_{\mathcal{X}}(g) \circ f$, will produce mn markers, where m and n are the number of markers produced by f and g respectively.

7 Conclusions and Related Work

We have adopted a labeled tree abstraction for a wide variety of structured and unstructured data sources, and we have shown that it is possible to develop simple query languages with an underlying optimizable algebra for querying and transforming labeled tree structures. We intend to incorporate at least the comprehension based fragment of UnQL into our existing implementation of CPL (Collection Programming Language [Won94, BDH⁺95]). However there are some important open questions.

Syntax. The syntax for making deep queries can probably be improved. There are deep restructuring operations expressible in UnCAL whose translation into *traverse ... giving ...* is rather cumbersome.

Expressive power. Certain restructuring queries appear not to be expressible in UnCAL because they require a fix-point operation. In view of Theorem 5.1 we know that transitive closure (expressed in terms of fixed-depth structures) is not expressible in UnQL. A more interesting problem is a generalization of the “edge merging” example at the end of Section 4. Merging sibling A edges at one level leads to the possibility that more A edges at the next level down will become siblings that can also be merged. It is easy to see that the result of repeatedly merging sibling A edges until no more merges are possible is well-defined, but it does not appear to be possible to compute this in UnQL.

Connections with “graph logic”. Another approach to querying labeled trees is to use some variety of “graph logic” [CM90]. It is possible to express at least the *vext* fragment of UnCAL in Datalog, given a graph representation of the input tree. We represent the tree by a set of $(oid, label, oid)$ triples and construct a program that expresses an output graph as a predicate on $(oid', label, oid')$ triples where oid' values are constructed by a Skolem function on the input *oids*, similar to id-terms in F-logic [KL89]. Whether Datalog opti-

mizations have anything to do with the optimizations we have presented here is an open question.

Other optimizations. Examination of the examples in Section 6 will indicate that, when a schema is present, one can obtain many useful optimizations from “pruning” the search, i.e., not searching subtrees that are known not to contain relevant structures. The problem here is how to describe schemas. We have seen how to do this for relational databases, and the same idea should generalize to object-oriented databases. However, there may be much looser notions of what a schema is for this labeled tree model.

References

- [BBW92] Val Breazu-Tannen, Peter Buneman, and Limsoon Wong. Naturally embedded query languages. In J. Biskup and R. Hull, editors, *LNCS 646: Proceedings of 4th International Conference on Database Theory, Berlin, Germany, October, 1992*, pages 140–154. Springer-Verlag, October 1992.
- [BDH⁺95] Peter Buneman, Susan Davidson, Kyle Hart, Chris Overton and Limsoon Wong. A Data Transformation System for Biological Data Sources. Proceedings of 21st VLDB, September 1995.
- [BDHS96] Peter Buneman, Susan Davidson, Gerd Hillebrand and Dan Suciu. A Query Language and Optimization Techniques for Unstructured Data. Technical Report MS-CIS-96-09, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104, February 1996.
- [BDS94] Peter Buneman, Susan Davidson, and Dan Suciu. Programming constructs for unstructured data. Technical Report MS-CIS-95-14, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104, March 1994.
- [BLS⁺94] P. Buneman, L. Libkin, D. Suciu, V. Tannen, and L. Wong. Comprehension syntax. *SIGMOD Record*, 23(1):87–96, March 1994.
- [BTS91] V. Breazu-Tannen and R. Subrahmanyam. Logical and computational aspects of programming with Sets/Bags/Lists. In *LNCS 510: Proceedings of 18th International Colloquium on Automata, Languages, and Programming, Madrid, Spain, July 1991*, pages 60–75. Springer Verlag, 1991.

- [CM90] M. Consens and A. Mendelzon. Graphlog: a visual formalism for real life recursion. In *Proceedings of the 9th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, April 1990.
- [GP84] A. Goldberg and R. Paige. Stream processing. In *Proceedings of ACM Symposium on LISP and Functional Programming*, pages 53–62, Austin, Texas, August 1984.
- [GRS93] N. Goodman, S. Rozen, and L. Stein. Requirements for a deductive query language in the MapBase genome-mapping database. In *Proceedings of Workshop on Programming with Logic Databases, Vancouver, BC*, October 1993.
- [KKS92] M. Kifer, W. Kim, and Y. Sagiv. Querying object-oriented databases. In M. Stonebraker, editor, *Proceedings ACM-SIGMOD International Conference on Management of Data*, pages 393–402, San Diego, California, June 1992.
- [KL89] M. Kifer and G. Lausen. F-logic: A higher order language for reasoning about objects, inheritance, and scheme. In *Proceedings of ACM-SIGMOD 1989*, pages 46–57, June 1989.
- [Kos95] Anthony S. Kosky. Observational Properties of Databases with Object Identity Technical Report MS-CIS-95-20, University of Pennsylvania, 1995.
- [Mil89] Robin Milner. *Communication and concurrency*. Prentice Hall, 1989.
- [PGMW95] Y. Papakonstantinou, H. Garcia-Molina, and J. Widom. Object exchange across heterogeneous information sources. In *IEEE International Conference on Data Engineering*, March 1995.
- [TMD92] J. Thierry-Mieg and R. Durbin. Syntactic Definitions for the ACEDB Data Base Manager. Technical Report MRC-LMB xx.92, MRC Laboratory for Molecular Biology, Cambridge, CB2 2QH, UK, 1992.
- [Won94] Limsoon Wong. *Querying Nested Collections*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104, August 1994.