

Purely Relational FLWORs

Torsten Grust
Technical University of Munich
Department of Computer Science, Database Systems
Munich, Germany
grust@in.tum.de

ABSTRACT

We report on a compilation procedure that derives relational algebra plans from arbitrarily nested XQuery FLWOR blocks. While recent research was able to develop relational encodings of trees which may turn RDBMSs into highly efficient XPath and XML Schema processors, here we describe *relational encodings of nested iteration, variables, and the item sequences* to which variables are bound. The developed techniques are *purely relational* in more than one sense: (a) we rely on a standard (or rather: classical) algebra that is readily supported by relational engines, and (b) we use relational concepts like functional and multivalued dependencies to significantly simplify the emitted plans. This work blends well with the mentioned tree encodings and thus contributes a further important building block to investigations into XQuery processors based on relational database technology.

1. INTRODUCTION

In a sense, the relational query processing model may be described as rather simplistic, on multiple levels: on the logical level, sequences of tuples (tables) are fed into algebraic operators which generate such sequences again, and the implementations of these operators work best if the tuple sequences are stored in a physically contiguous (secondary) memory fragment. While this simplicity is the key to the efficiency of relational query engines, it also prescribes a rather restrictive mode of data access and processing.

It is all the more remarkable that the database research community succeeded in turning RDBMSs into efficient processors for originally non-tabular *tree-shaped data*: trees are relationally encoded such that core operations on this data type, *e.g.*, path traversals or the validation of a tree's structure, are mapped onto sequential tuple scans which relational engines inherently support well.

Such tree encodings have primarily been used to devise relational XPath processors [9,11], but they may also yield efficient implementations of XML Schema validators [6]. Among

other benefits, the resulting systems inherit the scalability of the underlying relational back-ends [3]. It is legitimate to hope that this technology may be developed into full-fledged XQuery implementations, given that we can find relational ways to also express XQuery concepts beyond XPath axis traversals.

To this end, this paper does *not* talk about XPath evaluation at all but shifts focus to *the* central XQuery language feature, the **for-let-where-order-by-return** (or FLWOR) block [2]. The presence of arbitrarily nested iteration as well as the possibility to bind and then refer to variables in expressions presents an interesting technical challenge since the compilation target—relational algebra—is a combinator-style language that lacks any notion of explicit iteration or variables.

In Section 2, we start off by sketching the ideas which underly the compilation of nested XQuery FLWOR blocks into relational algebra. The canonical compilation yields relational plans featuring particularly simple instantiations of relational operators, *e.g.*, all occurring joins are equi-joins, but of significant size and certain redundancy: the item sequence representation is *denormalized* (Section 3). Due to the inherent simplicity of the plans, however, we are able to use a simple form of plan analysis to detect, among other properties, induced functional and multivalued dependencies which we then use in highly effective peep-hole-style plan simplification (Section 4). We will also sketch how the analysis of such dependencies may enable non-syntactic XQuery join detection (Section 5). The resulting relational XQuery processor supports the XQuery dialect sketched in Table 1 and is able to evaluate the XMark benchmark [12] queries on document sizes beyond 1 GB in interactive time (Section 6).

2. RELATIONAL FLWORs

We will now describe techniques which enable the compilation of the mentioned XQuery dialect (Table 1) into the relational algebra shown in Table 2. Besides the *row numbering* operator ϱ^1 , this constitutes a classical, rather restricted relational algebra variant. We will discuss the specifics of some operators as we go.

XQuery operates with *ordered, finite sequences of items* as the principal data type. *Items* either represent atomic values (of one of the XML Schema atomic types, *e.g.* **string**, **decimal**, **positiveInteger**) or nodes. For a relational representation of an atomic value v , we choose an implementa-

¹Many relational engines implement ϱ in terms of the DENSE_RANK operator defined by SQL:1999 [7].

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

XIME-P 2005, June 16–17, Baltimore, Maryland.

lifting occurs for the subexpressions $e_{1,2,3}$ of Q_1 have been marked by ② in Figure 1. From the loop-lifted e_1 we derive the representation $e_{\$x}$ of $\$x$ at ③. Next, the subexpressions occurring in the body of the outermost `for`-clause are to be compiled in iteration scope s_1 . The plan computes $loop_{s_1} = \pi_{iter}(e_{\$x})$ which is then used to loop-lift the relational representation of e_2 to yield the intermediate result at ④: this relation encodes the item sequence (30,20) for each of the three iterations performed by the outer `for`-clause. The compilation scheme continues, computes the representation of $\$y$ at ⑤ and emits algebraic code to derive the relational encoding ⑥ of the atomic constant 10 in the innermost iteration scope s_2 (via loop lifting, here: $loop_{s_2} \times \begin{smallmatrix} \text{pos} & \text{item} \\ 1 & 10 \end{smallmatrix}$): in each of the six iterations of the body of the inner `for`-clause, the constant assumes the value 10.

We really reap the benefits of this representation—which collects the values of a subexpression as well as the bindings of a variable *for all iterations* into single tables—when it comes to the bulk evaluation of XQuery functions and operators like, *e.g.*, `*` in Q_1 . At ⑦, we assemble the arguments $\$y$ and 10 to operator `*` in each iteration via an equi-join on the `iter` values. The subsequent \otimes operator then computes the overall six multiplications necessary to evaluate Q_1 *in bulk*. A similar remark applies to \ominus further up in the plan and to all operator and function applications in general.

The final result of the plan is ⑧: the fourth iteration of the innermost loop body contributes the single item 200, the fifth iteration contributes item 300 to the query result sequence (200,300).²

In a nutshell, loop lifting is the key technique to compile explicit iteration—XQuery’s `for`-clause—into efficient bulk-style applications of algebraic operators. We thus operate the relational engine in the mode it supports best.

3. PLAN SIZE AND DENORMALIZED ITEM SEQUENCES

It is a plus of this compilation scheme that it does not restrict the shape of the input query: expressions may nest arbitrarily given that the syntactical and typing rules of XQuery are obeyed. The emitted algebraic plans, however, may be large. Although a rather simple—and thus efficient—variant of relational algebra (Section 4) suffices to evaluate the plans, the significant number of operators may clearly have an impact on performance.

Note that the plans include several occurrences of common sub-plans such that a lot may already be gained by compiling into DAGs (instead of operator trees). The principal size problem remains, though. To exemplify, XMark query Q_8 (involving XPath navigation, a value-based join, aggregation, and element construction) compiles into a plan DAG of 153 operators. The simple query Q_2

```
for $x in (k, ..., 2, 1)
return $x * 5
```

(Q_2)

initially yields the plan of Figure 4(a) and thus a total of 13 operators.

²A final simple back-mapping step (not shown in Figure 1 for space reasons, but discussed in [8]) canonicalizes the values in columns `iter` and `pos`.

Denormalized item representation. Loop lifting effectively collects all values an expression e assumes during the evaluation of enclosing `for`-clauses into a single table. Should e be (a) *constant* or (b) *invariable* with respect to its enclosing iterations, loop lifting consequently leads to a *fully denormalized representation* for e and thus to—at least potentially—significant data redundancy: e ’s representation will be identical for all iterations (*i.e.*, groups of identical iter values).

Note the loop-lifted representation of the atomic item 5 in Q_2 shown in Figure 3(a). Similarly, inside a loop of k overall iterations, the constant or invariable item sequence (i_1, \dots, i_n) has the encoding depicted in Figure 3(b). Depending on the values of k and n , this redundancy may, again, severely impact plan performance.

XQuery’s orthogonality clearly is one of its virtues. We thus designed the compiler to be completely *compositional*—which is crucial and typical for a compiler of a functional-style language: a subexpression is compiled independent of its embedding expression. Plan size and data redundancy seem to be prices we pay for this orthogonal treatment of, *e.g.*, items and item sequences.

4. PLAN ANALYSIS AND SIMPLIFICATION

In the sequel, we argue that the just mentioned issues need not be showstoppers for loop-lifting-based XQuery compilation. The concluding Section 6 briefly reviews performance results which clearly indicate the same.

Despite their significant size, the generated plans come with an effective optimization hook: the operators are particularly simple, restricted variants of the classical relational algebra operators, *e.g.*,

- all \bowtie -operators are equi-joins,
- the projection operator π does not remove duplicates,
- all invocations of \cup receive disjoint argument relations.

These restrictions render the operators efficiently implementable inside a relational engine and—more importantly in this context—enable rather straightforward plan analysis to infer a variety of *properties of intermediate plan results*. We then use the properties to spot those places in the plans which are subject to simplification or even pruning. As we will see, the simplifications can effectively cut down plan size and help to remove inefficiencies due to denormalization through loop-lifting.

Because the relational plans have been generated by an XQuery compiler (as opposed to an SQL compiler) the simplification process has a somewhat different “taste” than regular algebraic optimization.

4.1 Plan Properties and Inference

Table 3 lists the properties we will use during plan analysis and simplification. Properties *icols* and *ocols* (*input/output columns*) are associated with operators, the remaining properties relate to intermediate plan results, *i.e.*, relations.

iter	pos	item
1	1	5
⋮	⋮	⋮
k	1	5

iter	pos	item
1	1	i_1
⋮	⋮	⋮
1	n	i_n
⋮	⋮	⋮
k	1	i_1
⋮	⋮	⋮
k	n	i_n

(a) (b)

Figure 3: Denormalized item sequences.

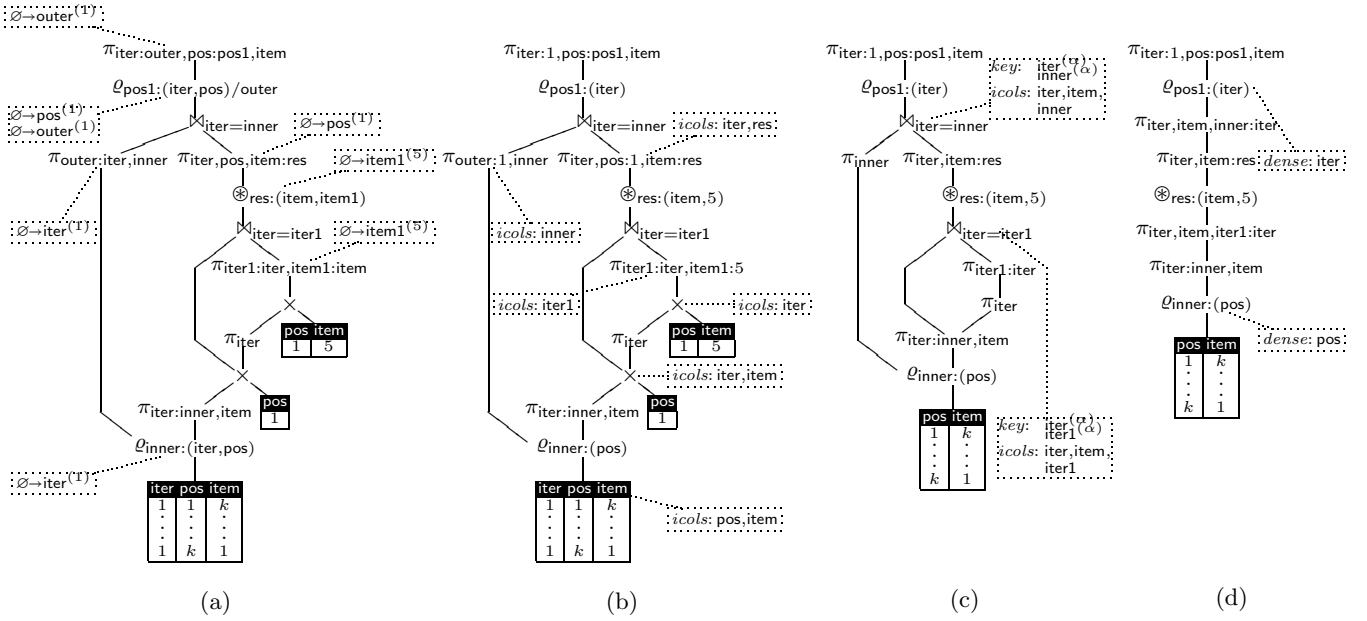


Figure 4: Plan simplification based on inferred properties of intermediate results and operators (query Q_2).

Property	Description
$card_0, card_1$	relation has cardinality 0 (or 1)
$key: a^{(\alpha)}$	column a is key with value domain α
$dense: a$	col. a is numbered densely starting from 1
$\emptyset \rightarrow a^{(v)}$	a is constant col. of value v
$\emptyset \rightarrow a_1^{(\alpha_1)}, \dots, a_n^{(\alpha_n)}$	remaining cols. are independent of a_1, \dots, a_n
$icols: a_1, \dots, a_n$	cols. a_1, \dots, a_n are required to evaluate this operator (or an ancestor operator)
$ocols: a_1^{(\alpha_1)}, \dots, a_n^{(\alpha_n)}$	cols. a_1, \dots, a_n appear in operator output

Table 3: Properties of relations (and operators).

The two former properties are primarily used to detect that operators consume or produce obsolete result columns—a situation that is commonly created by preceding simplification steps. An operator ω inherits property $icols$ from its parents in the plan DAG: if column a is required to evaluate a parent and ω does not produce output column a itself, then $\omega.icols$ will contain column a (and ω thus requires this column to be contained in one of its arguments). If, for example, $\omega = \pi_{a:b,c}$ or $\omega = \otimes_{a:(b,c)}$, $\omega.icols$ would contain columns b, c instead (but not column a).

All other properties may be synthesized during a single bottom-up walk of the DAG. An excerpt of the property inference rules can be found in Figure 5. Property $e.card_0$ ($e.card_1$) holds if sub-plan e yields a relation of exactly zero (one) tuple(s). From this we may infer further interesting relation characteristics, *e.g.*, the preservation of keys (see below). The presence of such a key column a is recorded in property $key: a^{(\alpha)}$. A superscript (α) denotes a *domain identifier*: α represents the active domain of its column, *i.e.*, all values actually contained in column a . We obviously cannot maintain the value domain itself due to its potential size, but two columns $a^{(\alpha)}, b^{(\alpha)}$ with identical inferred domain identifiers are guaranteed to share one active domain. A device similar to domain identifiers has also been used in [10]. In a relation e with property $e.dense: a$, the active domain

of column a is $1, 2, \dots, n$ where n is the cardinality of e .

The two remaining properties are probably the most important ones in our context. If $e.\emptyset \rightarrow a^{(v)}$, then all tuples in e contain value v in column a . We denote this by a degenerate *functional dependency* (fd) with an empty left-hand side [1, 5, 10] for uniformity reasons which will become clear in a moment. The degenerate *multivalued dependency* (mvd) $e.\emptyset \rightarrow a_1^{(\alpha_1)}, \dots, a_n^{(\alpha_n)}$ indicates that the column group a_1, \dots, a_n is independent of all remaining columns of e . For the relation in Figure 3(a), for example, we have $\emptyset \rightarrow item^{(5)}$. For the relation in Figure 3(b), we have $\emptyset \rightarrow iter^{(\alpha)}$ where α represents the domain $1, \dots, k$.

In Figure 5, Rule (5) uses inferred functional dependencies for constant folding with respect to the binary operator \circ , while Rule (1) introduces such fds. Result column a will be constant in both cases. Rule (2) propagates key column b through a selection but introduces a fresh domain identifier β for b : the selection potentially restricts active domain α ; we know, however, that $\beta \subseteq \alpha$. This domain inclusion relationship is recorded for later perusal. Finally, an mvd is inferred by Rule (9). The Cartesian product guarantees the column values in e_1 and e_2 to be independent or orthogonal [5]. (Obviously, a symmetric inference rule in which $e_{1,2}$ swap roles is also correct.) Note that such mvds are preserved under certain preconditions (Rules (10)–(12)).

4.2 Peep-Hole Plan Simplification

The substantial size of the initially emitted plans make a classical approach to query rewriting a rather hopeless case: rewrite rules would need to be either (a) very specific, and thus carry huge DAG contexts, or (b) very generic which makes their applicability hard to detect. Instead, we use a collection of *peep-hole*-style equivalences (Figure 6) whose applicability may be decided by looking at the inferred properties of a *single* plan DAG node. Specific necessary properties thus appear in the premises of the algebraic equivalence rules.

$$\begin{array}{c}
\frac{}{(\sigma_a(e)).\varnothing \rightarrow \mathbf{a}^{(true)}} \text{ (1)} \quad \frac{e.key: \mathbf{b}^{(\alpha)}}{(\sigma_a(e)).key: \mathbf{b}^{(\beta)} \wedge \beta \subseteq \alpha} \text{ (2)} \quad \frac{}{(\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})}(e)).dense: \mathbf{a} \wedge (\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})}(e)).key: \mathbf{a}^{(\alpha)}} \text{ (3)} \quad \frac{e.\varnothing \rightarrow \mathbf{d}^{(\alpha)}}{(\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})/d}(e)).dense: \mathbf{a} \wedge (\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})/d}(e)).key: \mathbf{a}^{(\beta)}} \text{ (4)} \\
\frac{e.\varnothing \rightarrow \mathbf{b}^{(v_1)} \quad e.\varnothing \rightarrow \mathbf{c}^{(v_2)}}{(\varrho_{\mathbf{a}: (\mathbf{b}, \mathbf{c})}(e)).\varnothing \rightarrow \mathbf{a}^{(v_1 \vee v_2)}} \text{ (5)} \quad \frac{e_1.\varnothing \rightarrow \mathbf{a}^{(v)} \quad e_2.\varnothing \rightarrow \mathbf{a}^{(v)}}{(e_1 \dot{\cup} e_2).\varnothing \rightarrow \mathbf{a}^{(v)}} \text{ (6)} \quad \frac{e_1.card_1 \quad e_2.key: \mathbf{a}^{(\alpha)}}{(e_1 \times e_2).key: \mathbf{a}^{(\alpha)}} \text{ (7)} \quad \frac{e_1.key: \mathbf{a}^{(\alpha)} \quad e_2.key: \mathbf{b}^{(\beta)} \quad \beta \subseteq \alpha}{(e_1 \bowtie_{\mathbf{a}=\mathbf{b}} e_2).key: \mathbf{a}^{(\beta)}, \mathbf{b}^{(\beta)}} \text{ (8)} \\
\frac{e_2.ocols: \mathbf{a}_1^{(\alpha_1)}, \dots, \mathbf{a}_n^{(\alpha_n)}}{(e_1 \times e_2).\varnothing \rightarrow \mathbf{a}_1^{(\alpha_1)}, \dots, \mathbf{a}_n^{(\alpha_n)}} \text{ (9)} \quad \frac{e_1.\varnothing \rightarrow \mathbf{a}^{(\alpha)} \quad e_2.\varnothing \rightarrow \mathbf{a}^{(\alpha)}}{(e_1 \dot{\cup} e_2).\varnothing \rightarrow \mathbf{a}^{(\alpha)}} \text{ (10)} \quad \frac{e.\varnothing \rightarrow \mathbf{a}_1^{(\alpha_1)}, \dots, \mathbf{a}_n^{(\alpha_n)}}{(\pi_{\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{b}, \dots, \mathbf{c}}(e)).\varnothing \rightarrow \mathbf{a}_1^{(\alpha_1)}, \dots, \mathbf{a}_n^{(\alpha_n)}} \text{ (11)} \quad \frac{e.\varnothing \rightarrow \mathbf{d}^{(\alpha)}}{(\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})/d}(e)).\varnothing \rightarrow \mathbf{d}^{(\alpha)}} \text{ (12)}
\end{array}$$

Figure 5: Property inference (excerpt of rules). Domain identifiers α, β introduced in a rule’s consequence denote fresh (previously unused) domain identifiers.

$$\begin{array}{c}
\frac{e.\varnothing \rightarrow \mathbf{b}^{(v)}}{\pi_{\mathbf{a}, \dots, \mathbf{b}, \dots, \mathbf{c}}(e) \equiv \pi_{\mathbf{a}, \dots, \mathbf{b}:v, \dots, \mathbf{c}}(e)} \text{ (a)} \quad \frac{e.\varnothing \rightarrow \mathbf{b}^{(v)}}{\varrho_{\mathbf{a}: (\mathbf{b}, \mathbf{c})}(e) \equiv \varrho_{\mathbf{a}: (v, \mathbf{c})}(e)} \text{ (b)} \quad \frac{e.\varnothing \rightarrow \mathbf{d}^{(v)}}{\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})/d}(e) \equiv \varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})}(e) \wedge \varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{d}, \dots, \mathbf{c})/e}(e) \equiv \varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})/e}(e)} \text{ (c)} \\
\frac{e.key: \mathbf{d}^{(\alpha)}}{\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{d}, \dots, \mathbf{c})/e}(e) \equiv \varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{d})/e}(e)} \text{ (d)} \quad \frac{e.dense: \mathbf{b} \quad (\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})}(e)).icols: \mathbf{d}_1, \dots, \mathbf{d}_n}{\varrho_{\mathbf{a}: (\mathbf{b}, \dots, \mathbf{c})}(e) \equiv \pi_{\mathbf{a}: \mathbf{b}, \mathbf{d}_1, \dots, \mathbf{d}_n}(e)} \text{ (e)} \quad \frac{(e_1 \times e_2).icols \subseteq e_2.ocols \quad e_1.card_1}{e_1 \times e_2 \equiv e_2} \text{ (f)} \\
\frac{e_1.key: \mathbf{a}^{(\alpha)} \quad e_2.key: \mathbf{b}^{(\beta)} \quad \beta \subseteq \alpha \quad (e_1 \bowtie_{\mathbf{a}=\mathbf{b}} e_2).icols: \mathbf{d}_1, \dots, \mathbf{d}_n \subseteq e_2.ocols}{e_1 \bowtie_{\mathbf{a}=\mathbf{b}} e_2 \equiv \pi_{\mathbf{a}: \mathbf{b}, \mathbf{d}_1, \dots, \mathbf{d}_n}(e_2)} \text{ (g)} \quad \frac{e_1.key: \mathbf{a}^{(\alpha)} \quad e_2.\varnothing \rightarrow \mathbf{b}^{(\alpha)}}{e_1 \bowtie_{\mathbf{a}=\mathbf{b}} e_2 \equiv e_1 \times e_2} \text{ (h)}
\end{array}$$

Figure 6: Algebraic equivalences (excerpt) based on inferred properties.

Some simplifications strive to minimize the dependence of an operator on specific input columns (in Rules (a) and (b), a reference to constant column \mathbf{b} is traded for the constant value v). This removes column \mathbf{b} from the *icols* property which, in turn, may render a Cartesian product (Rule (f)) or a join (Rule (g)) further down in the DAG obsolete. Another group of equivalences tries to simplify or prune occurrences of ϱ : the semantics of ϱ are simple but most conceivable implementations require its input to be sorted. Rules (c) and (d) describe how a constant or key column \mathbf{d} may be used to shorten the list of sort criteria. In Rule (d), column \mathbf{d} is a key for e and thus a decisive order criterion: further columns need not be inspected to decide the order of two rows. If a dense column \mathbf{b} is the major sort criterion, we might as well derive the row numbering from \mathbf{b} itself; ϱ is pruned (Rule (e)). Rule (g) finally removes key joins in case of a key domain inclusion: every row in e_2 is guaranteed to find exactly one join partner in e_1 . Note that we do not carry out the actual domain inclusion test here but rather check domain identifiers against the known inclusion relationships recorded by inference Rule (2).

Figure 4 exemplifies how these rules interact to simplify the initial plan for query Q_2 in several steps. Those nodes affected by the following simplification step are annotated with the properties which justify the simplification (remember that we are using peep-hole simplification: changes to the DAG are always in the direct vicinity of the affected nodes).

The bottom ϱ node in Figure 4(a), for example, is simplified according to Rule (c): we have inferred that all rows carry the atomic value 1 in column *item* ($\varnothing \rightarrow \text{item}^{(1)}$). Such a column does not make for a useful order criterion and is




Figure 7: Fully simplified plan for query Q_2 .


removed from the list of order criteria. The resulting, simpler row number operator $\varrho_{\text{inner}: (\text{pos})}$ numbers the incoming rows solely based on the dense column *pos* (Figure 4(d)). In such a situation, Rule (e) indicates that the dense column already provides the requested row numbering *for free*. This instance of ϱ is thus replaced by a projection $\pi_{\text{inner}: \text{pos}}$ which introduces column *inner* as a mere alias for column *pos*. In Figure 7, which depicts the fully simplified plan for query Q_2 , even this projection has been pruned away.

Similarly, note how the loop lifting for the atomic item 5 is detected (and then rendered obsolete) with the help of property $\varnothing \rightarrow \text{item}^{(5)}$ in the first two simplification steps in Figure 4.

5. MVDs FOR XQUERY JOIN DETECTION

An $\text{fd } \varnothing \rightarrow \mathbf{a}^{(v)}$ can detect the presence of a loop-lifted constant item in the plan DAGs. Likewise—and we feel that this is a quite elegant outcome of this research—we may use the more general mvd concept to detect the more general case of a loop-lifted invariable *item sequence*. To see this, note that the representation of the loop-lifted item sequence (i_1, \dots, i_n) (Figure 3(b)) is derived from the relational sequence encoding (see Section 2) by forming the Cartesian product with the current *loop* relation, *i.e.*, with a single-column relation with schema *iter*. But this *precisely* characterizes [1, 5] the presence of the degenerate mvd $\varnothing \rightarrow \text{iter}^{(\alpha)}$ (where α is determined by the number of iterations to be performed) in the encoding of Figure 3(b). The property inference Rule (9) in Figure 5 makes this correspondence between loop-lifting and the presence of mvds available to the simplifier.

This observation paves the way for a form of XQuery join detection that is completely unaffected by XQuery’s syntactic diversity. Figure 8(a) shows a prototypical plan in which the (arbitrary) subexpression e_2 has been lifted over the iteration scope opened by subexpression e_1 (which itself is to



(a) Initial plan: sub-plan e_2 loop-lifted over e_1 . (b) Equivalent plan: sub-plan e_2 not loop-lifted over e_1 .

Figure 8: An inferred mvd (at \ominus) helps to recognize an XQuery join scenario.




Figure 9: Fully simplified plan for query Q_1 .

be evaluated inside an iteration encoded by relation *loop*). Exactly if the source XQuery query describes a join between the two subexpressions $e_{1,2}$ —*i.e.*, e_2 may be evaluated independently of e_1 before the two resulting item sequences are joined in some way—then we will be able to infer the mvd $\emptyset \rightarrow \text{iter}^{(\alpha)}$ (marked \ominus in Figure 8(a)) at the root of the sub-plan for e_2 . We then simplify with the help of Rule (h) in Figure 6 and emit the plan of Figure 8(b) instead. In the simpler plan, note that e_2 is *not* loop-lifted over e_1 but over *loop*. We may now expect (the larger the intermediate result of e_1 the more significant) less data redundancy in e_2 . More importantly, however, the new plan is nicely prepared to merge its root Cartesian product with a σ -operator to form a join. This algebraic join then implements the join that was present in the original XQuery query in some syntactic form.

The initial plan (Figure 1) for query Q_1 is an instance of the prototypical plan of Figure 8(a). Note that the compiler can detect the presence of the mvd $\emptyset \rightarrow \text{iter}^{(\alpha)}$ at node \ominus in Figure 1 *while* it emits the code. At this point in time, the compiler holds references to the sub-plans for $e_{1,2}$ and can rather easily initiate the rewriting into the new DAG shape of Figure 8(b). The newly introduced Cartesian product is then merged with the two ancestor operators σ - \ominus to make the implicit value-based join expressed by XQuery query Q_1 explicit in the algebraic plan. Note that, depending on the XQuery comparison operator used in the source query, this join operator will be a θ -join (not an equi-join). In [3], we discuss how the existential semantics of XQuery’s general comparison operators ($=$, $<$, ...) may be compiled into efficient algebraic code. Figure 9 shows the fully simplified

plan for query Q_1 . This variant of XQuery join detection is not affected by the shape of the source query. We will, in particular, infer the mvd $\emptyset \rightarrow \text{iter}^{(\alpha)}$ at \ominus in Figure 8(a) independent of how many iteration scopes, *i.e.*, nested *for*-clauses, separate e_1 and e_2 .

6. PURELY RELATIONAL XQUERY

The XQuery compiler studied in this paper, dubbed *Pathfinder*, is part of the purely relational XQuery processor *MonetDB/XQuery* [3]. The back-end of this processor is *MonetDB* [4], an extensible open-source relational database kernel developed at CWI, Amsterdam. MonetDB’s internals have been optimized for in-memory operation but the system is engineered to exploit the memory and file management services of the operating system to operate with data volumes that exceed main memory capacity. In particular, MonetDB efficiently supports *narrow tables* (three columns or less) which are pervasive in the plans emitted by our XQuery compiler.

To give an impression of where research into relational XQuery processors may lead, we reproduce a fragment of the XMark experiments reported in [3] in Table 4 (conducted on a 64-bit system with an 1.6 GHz AMD processor and 8 GB RAM). Pathfinder, faithfully operating MonetDB in the strict sequential access mode it was engineered for, indeed inherits the scalability of its back-end: XMark document instances beyond 1GB serialized size are processed in interactive time, including the XQuery join queries Q_8 – Q_{10} (note that queries Q_{11} , Q_{12} produce intermediate results quadratic in the size of the input document—any XQuery processor will thus struggle at this point).

Pathfinder and the MonetDB/XQuery system will be released into the open-source well before XIME-P 2005 (www.pathfinder-xquery.org).

Q	110 MB	1.1 GB
1	0.41	1.2
2	0.30	2.4
3	1.51	12.5
4	0.45	3.8
5	0.16	1.2
6	0.05	0.3
7	0.07	0.4
8	0.75	10.4
9	0.87	12.9
10	5.31	55.0
11	3.48	960.9
12	1.66	431.3
13	0.22	1.3
14	2.20	21.3
15	0.28	1.7
16	0.26	1.8
17	0.30	2.8
18	0.13	0.9
19	0.55	5.3
20	0.62	4.9

Table 4: XMark query evaluation time (seconds).

7. REFERENCES

- [1] C. Beeri, R. Fagin, and J.H. Howard. A Complete Axiomatization for Functional and Multivalued Dependencies in Database Relations. In *Proc. of the Int’l ACM SIGMOD Conference on Management of Data*, Toronto, Canada, August 1977.
- [2] S. Boag, D. Chamberlin, M.F. Fernández, D. Florescu, J. Robie, and J. Siméon. XQuery 1.0: An XML Query Language. World Wide Web Consortium, February 2005. W3C Working Draft <http://www.w3.org/TR/xquery/>.
- [3] P. Boncz, T. Grust, S. Manegold, J. Rittinger, and J. Teubner. Pathfinder: Relational XQuery Over Multi-Gigabyte Inputs in Interactive Time. Technical Report INS-E0503, CWI, Amsterdam, March 2005.
- [4] P. Boncz and M.L. Kersten. MIL Primitives for Querying a Fragmented World. *VLDB Journal*, 8(2), March 1999.
- [5] R. Fagin. Multivalued Dependencies and a New Normal Form for Relational Databases. *ACM Transactions on Database Systems (TODS)*, 2(3), September 1977.
- [6] T. Grust and S. Klinger. Validation and Type Annotation for Encoded Trees. In *Proc. of the 1st Int’l Workshop on XQuery Implementation, Experience and Perspectives (XIME-P)*, Paris, France, June 2004.

- [7] T. Grust, S. Sakr, and J. Teubner. XQuery on SQL Hosts. In *Proc. of the 30th Int'l Conference on Very Large Databases (VLDB)*, Toronto, Canada, September 2004.
- [8] T. Grust and J. Teubner. Relational Algebra: Mother Tongue—XQuery: Fluent. In *Proc. of the Twente Data Management Workshop on XML Databases and Information Retrieval (TDM)*, The Netherlands, June 2004.
- [9] T. Grust, M. van Keulen, and J. Teubner. Staircase Join: Teach a Relational DBMS to Watch Its Axis Steps. In *Proc. of the 29th Int'l Conference on Very Large Databases (VLDB)*, Berlin, Germany, September 2003.
- [10] A. Klug. Calculating Constraints on Relational Expressions. *ACM Transactions on Database Systems (TODS)*, 5(3), September 1980.
- [11] P. O'Neil, E. O'Neil, S. Pal, I. Cseri, G. Schaller, and N. Westbury. ORDPATHs: Insert-Friendly XML Node Labels. In *Proc. of the Int'l ACM SIGMOD Conference on Management of Data*, Paris, France, June 2004.
- [12] A. Schmidt, F. Waas, M.L. Kersten, M.J. Carey, I. Manolescu, and R. Busse. XMark: A Benchmark for XML Data Management. In *Proc. of the 28th Int'l Conference on Very Large Databases (VLDB)*, Hong Kong, China, August 2002.