

Flattening Nested Database Queries

Ben Lippmeier

University of New South Wales

FP-SYD 2014/05/28

Nested array representation

`arr` = `[[1 2] [3 4 5] [] [6]]`

`segd` = `[2 3 0 1]`

`data` = `[1 2 3 4 5 6]`

Flat vs Nested Data Parallelism

- Flat Parallelism: Worker function is sequential.

```
thingo xs
  = mapP (\x. x + 1) xs
```

- Nested Parallelism: Worker function is parallel.

```
thingo xss
  = mapP (\xs. zipWithP g xs ys) xss
```

- The Flattening / Vectorisation transform converts nested parallelism into flat parallelism.

The Flattening Transform

```
f :: Int -> Int
```

```
f x = x + 1
```

```
g :: Array Int -> Array Int
```

```
g ys = mapP f ys
```

The Flattening Transform

```
f :: Int -> Int
```

```
f x = x + 1
```

```
g :: Array Int -> Array Int
```

```
g ys = mapP f ys
```

```
fL :: Array Int -> Array Int
```

```
fL xs = xs +L (replicate n 1)
```

```
  where n = length xs
```

```
g :: Array Int -> Array Int
```

```
g ys = fL ys
```

The Flattening Transform

```
f :: Int -> Int
```

```
f x = x + 1
```

```
g :: Array Int -> Array Int
```

```
g ys = mapP f ys
```

```
fL :: Array Int -> Array Int
```

```
fL xs = zipWithP (+) xs (replicate n 1)
```

```
  where n = length xs
```

```
g :: Array Int -> Array Int
```

```
g ys = fL ys
```

The Flattening Transform

```
f :: Int -> Int
```

```
f x = x + 1
```

```
g :: Array Int -> Array Int
```

```
g ys = mapP f ys
```

```
h :: Array (Array Int) -> Array (Array Int)
```

```
h zss = mapP g zss
```

The Flattening Transform

```
f :: Int -> Int
```

```
f x = x + 1
```

```
g :: Array Int -> Array Int
```

```
g ys = mapP f ys
```

```
h :: Array (Array Int) -> Array (Array Int)
```

```
h zss = mapP g zss
```

```
g :: Array Int -> Array Int
```

```
g ys = fL ys
```

```
gL :: Array (Array Int) -> Array (Array Int)
```

```
gL yss = fLL yss
```


The Flattening Transform

```
f :: Int -> Int
```

```
f x = x + 1
```

```
g :: Array Int -> Array Int
```

```
g ys = mapP f ys
```

```
h :: Array (Array Int) -> Array (Array Int)
```

```
h zss = mapP g zss
```

```
g :: Array Int -> Array Int
```

```
g ys = fL ys
```

```
gL :: Array (Array Int) -> Array (Array Int)
```

```
gL yss = unconcatP yss (fL (concatP yss))
```

Nested array representation

`arr` = `[[1 2] [3 4 5] [] [6]]`

`segd` = `[2 3 0 1]`

`data` = `[1 2 3 4 5 1]`

Relational Algebra: Mother Tongue—XQuery: Fluent

Torsten Grust Jens Teubner
University of Konstanz
Department of Computer & Information Science
Box D 188, 78457 Konstanz, Germany
{grust,teubner}@inf.uni-konstanz.de

ABSTRACT

This work may be seen as a further proof of the versatility of the relational database model. Here, we add XQuery to the catalog of languages which RDBMSs are able to “speak” fluently.

Given suitable relational encodings of sequences and ordered, unranked trees—the two data structures that form the backbone of the XML and XQuery data models—we describe a compiler that translates XQuery expressions into a simple and quite standard relational algebra which we expect to be efficiently implementable on top of any relational query engine. The compilation procedure is fully compositional and emits algebraic code that strictly adheres to the XQuery language semantics: document and sequence order as well as node identity are obeyed. We exercise special care in translating arbitrarily nested XQuery `FLWOR` iteration constructs into equi-joins, an operation which RDBMSs can perform particularly fast. The resulting purely relational XQuery processor shows promising performance figures in experiments.

types onto tables. Such encodings have also been proposed for *ordered, unranked trees*, the data type that forms the backbone of the XML data model. These mappings turn RDBMSs into *relational XML processors*. Furthermore, if the tree encoding is designed such that core operations on trees—XPath axis traversals—lead to efficient table operations, this can result in high-performance *relational XPath* implementations [8, 10].

In this work we extend the relational XML processing stack and propose the fully relational evaluation of XQuery [1] expressions. We give a compositional set of translation rules that compile XQuery expressions into a standard, quite primitive relational algebra. We expect any relational query engine to be able to efficiently implement the operators of this algebra. The operators were, in fact, designed to match the capabilities of modern SQL-based relational database systems (*e.g.*, the row numbering operator ρ exactly mirrors SQL:1999 OLAP ranking functionality) [9].

By design, we only have minimalistic assumptions on the underlying tree encoding, met by several XML encoding schemes [4, 13]. Our algebra can be easily modified to over-

Nested XQueries

$$s \left\{ \begin{array}{l} \text{for } \$v_0 \text{ in } (1,2) \text{ return} \\ \quad s_0 \left\{ \begin{array}{l} (\$v_0, \\ \quad \text{for } \$v_{0.0} \text{ in } (10,20) \text{ return} \\ \quad \quad s_{0.0} \{ (\$v_0, \$v_{0.0}) \} \\ \quad \quad) \end{array} \right. \end{array} \right.$$

Nested XQueries

$$s \left\{ \begin{array}{l} \text{for } \$v_0 \text{ in } (1,2) \text{ return} \\ \quad s_0 \left\{ \begin{array}{l} (\$v_0, \\ \text{for } \$v_{0.0} \text{ in } (10,20) \text{ return} \\ \quad s_{0.0} \{ (\$v_0, \$v_{0.0}) \} \end{array} \right. \\ \end{array} \right.$$

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"10"
2	1	"1"
2	2	"20"
3	1	"2"
3	2	"10"
4	1	"2"
4	2	"20"

(a) Intermediate result in $s_{0.0}$.

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"10"
1	3	"1"
1	4	"20"
2	1	"2"
2	2	"10"
2	3	"2"
2	4	"20"

(b) Intermediate result in s_0 .

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"1"
1	3	"10"
1	4	"1"
1	5	"20"
1	6	"2"
1	7	"2"
1	8	"10"
1	9	"2"
1	10	"20"

(c) Final result in top-level scope.

Nested array representation

`arr` = `[[1 2] [3 4 5] [] [6]]`

`segd` = `[2 3 0 1]`

`data` = `[1 2 3 4 5 1]`

Nested array representation

arr = [[1 2] [3 4 5] [] [6]]

segd = [2 3 0 1]

data = [1 2 3 4 5 1]

segdA' = [0 0 1 1 1 3]

segdB' = [0 1 0 1 2 0]

$e ::= c$	atomic constants
$\$v$	variables
(e, e)	sequence construction
$e/\alpha :: n$	loc. step (axis α , node test n)
element $t \{ e \}$	element constructor (tag t)
for $\$v$ in e return e	iteration
let $\$v := e$ return e	let binding
$e+e$	addition

$e ::= c$	atomic constants
$\$v$	variables
(e, e)	sequence construction
$e/\alpha::n$	loc. step (axis α , node test n)
$\text{element } t \{ e \}$	element constructor (tag t)
$\text{for } \$v \text{ in } e \text{ return } e$	iteration
$\text{let } \$v := e \text{ return } e$	let binding
$e+e$	addition

$\pi_{a_1:b_1, \dots, a_n:b_n}$	projection (and renaming)
σ_a	selection
$\dot{\cup}$	disjoint union
\times	cartesian product
$\bowtie_{a=b}$	equi-join
$\rho_{b:\langle a_1, \dots, a_n \rangle / p}$	row numbering
$\lceil \alpha, n$	XPath axis join (axis α , node test n)
ε	element construction
$\textcircled{*}_{b:\langle a_1, \dots, a_n \rangle}$	n -ary arithmetic/comparison operator $*$
$\underline{a b}$	literal table

$$\Gamma; \text{loop}; \Delta \vdash e \Rightarrow (q, \Delta')$$

$$\frac{}{\Gamma; \text{loop}; \Delta \vdash c \Rightarrow \left(\text{loop} \times \frac{\text{pos} \mid \text{item}}{1 \mid c}, \Delta \right)}. \quad (\text{CONST})$$

$$\mathbb{L}_n[c] = \text{replicate } n \ c$$

$$\Gamma; \text{loop}; \Delta \vdash e \Rightarrow (q, \Delta')$$

$$\Gamma; \text{loop}; \Delta \vdash e_1 \Rightarrow (q_1, \Delta_1) \quad \Gamma; \text{loop}; \Delta_1 \vdash e_2 \Rightarrow (q_2, \Delta_2)$$

$$\Gamma; \text{loop}; \Delta \vdash e_1 + e_2 \Rightarrow (\pi_{\text{iter}, \text{pos}, \text{item}: \text{res}} (\oplus_{\text{res}: \langle \text{item}, \text{item}' \rangle} (q_1 \bowtie_{\text{iter}=\text{iter}'} (\pi_{\text{iter}': \text{iter}, \text{item}': \text{item}} q_2))), \Delta_2)$$

$$\Gamma; \text{loop}; \Delta \vdash e \Rightarrow (q, \Delta')$$

$$\Gamma; \text{loop}; \Delta \vdash e_1 \Rightarrow (q_1, \Delta_1) \quad \Gamma; \text{loop}; \Delta_1 \vdash e_2 \Rightarrow (q_2, \Delta_2)$$

$$\Gamma; \text{loop}; \Delta \vdash e_1 + e_2 \Rightarrow (\pi_{iter, pos, item:res} (\oplus_{res:\langle item, item' \rangle} (q_1 \bowtie_{iter=iter'} (\pi_{iter':iter, item':item} q_2))), \Delta_2)$$

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"10"
2	1	"1"
2	2	"20"
3	1	"2"
3	2	"10"
4	1	"2"
4	2	"20"

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"10"
1	3	"1"
1	4	"20"
2	1	"2"
2	2	"10"
2	3	"2"
2	4	"20"

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"1"
1	3	"10"
1	4	"1"
1	5	"20"
1	6	"2"
1	7	"2"
1	8	"10"
1	9	"2"
1	10	"20"

(a) Intermediate result in $s_{0.0}$.

(b) Intermediate result in s_0 .

(c) Final result in top-level scope.

$$\Gamma; \text{loop}; \Delta \vdash e \Rightarrow (q, \Delta')$$

$$\Gamma; \text{loop}; \Delta \vdash e_1 \Rightarrow (q_1, \Delta_1) \quad \Gamma; \text{loop}; \Delta_1 \vdash e_2 \Rightarrow (q_2, \Delta_2)$$

$$\Gamma; \text{loop}; \Delta \vdash e_1 + e_2 \Rightarrow (\pi_{iter, pos, item:res} (\oplus_{res:\langle item, item' \rangle} (q_1 \bowtie_{iter=iter'} (\pi_{iter':iter, item':item} q_2))), \Delta_2)$$

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"10"
2	1	"1"
2	2	"20"
3	1	"2"
3	2	"10"
4	1	"2"
4	2	"20"

(a) Intermediate result in $s_{0.0}$.

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"10"
1	3	"1"
1	4	"20"
2	1	"2"
2	2	"10"
2	3	"2"
2	4	"20"

(b) Intermediate result in s_0 .

<i>iter</i>	<i>pos</i>	<i>item</i>
1	1	"1"
1	2	"1"
1	3	"10"
1	4	"1"
1	5	"20"
1	6	"2"
1	7	"2"
1	8	"10"
1	9	"2"
1	10	"20"

(c) Final result in top-level scope.

$$\mathbb{L}_n [e_1 + e_2] = \mathbb{L}_n [e_1] + \mathbb{L} \mathbb{L}_n [e_2]$$

$$\begin{array}{l}
\{\dots, \$v_i \mapsto q_{v_i}, \dots\}; \text{loop}; \Delta \vdash e_1 \Rightarrow (q_1, \Delta_1) \quad q_v \equiv \frac{pos}{I} \times \pi_{iter:inner,item} (\varrho_{inner:\langle iter,pos \rangle} q_1) \\
\text{loop}_v \equiv \pi_{iter} q_v \quad \text{map} \equiv \pi_{outer:iter,inner} (\varrho_{inner:\langle iter,pos \rangle} q_1) \\
\frac{\{\dots, \$v_i \mapsto \pi_{iter:inner,pos,item} (q_{v_i} \bowtie_{iter=outer} \text{map}), \dots\} + \{\$v \mapsto q_v\}; \text{loop}_v; \Delta_1 \vdash e_2 \Rightarrow (q_2, \Delta_2)}{\{\dots, \$v_i \mapsto q_{v_i}, \dots\}; \text{loop}; \Delta \vdash \text{for } \$v \text{ in } e_1 \text{ return } e_2 \Rightarrow} \\
(\pi_{iter:outer,pos:pos_1,item} (\varrho_{pos_1:\langle iter,pos \rangle/outer} (q_2 \bowtie_{iter=inner} \text{map})), \Delta_2)
\end{array}$$

(FOR)

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 re-

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

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.

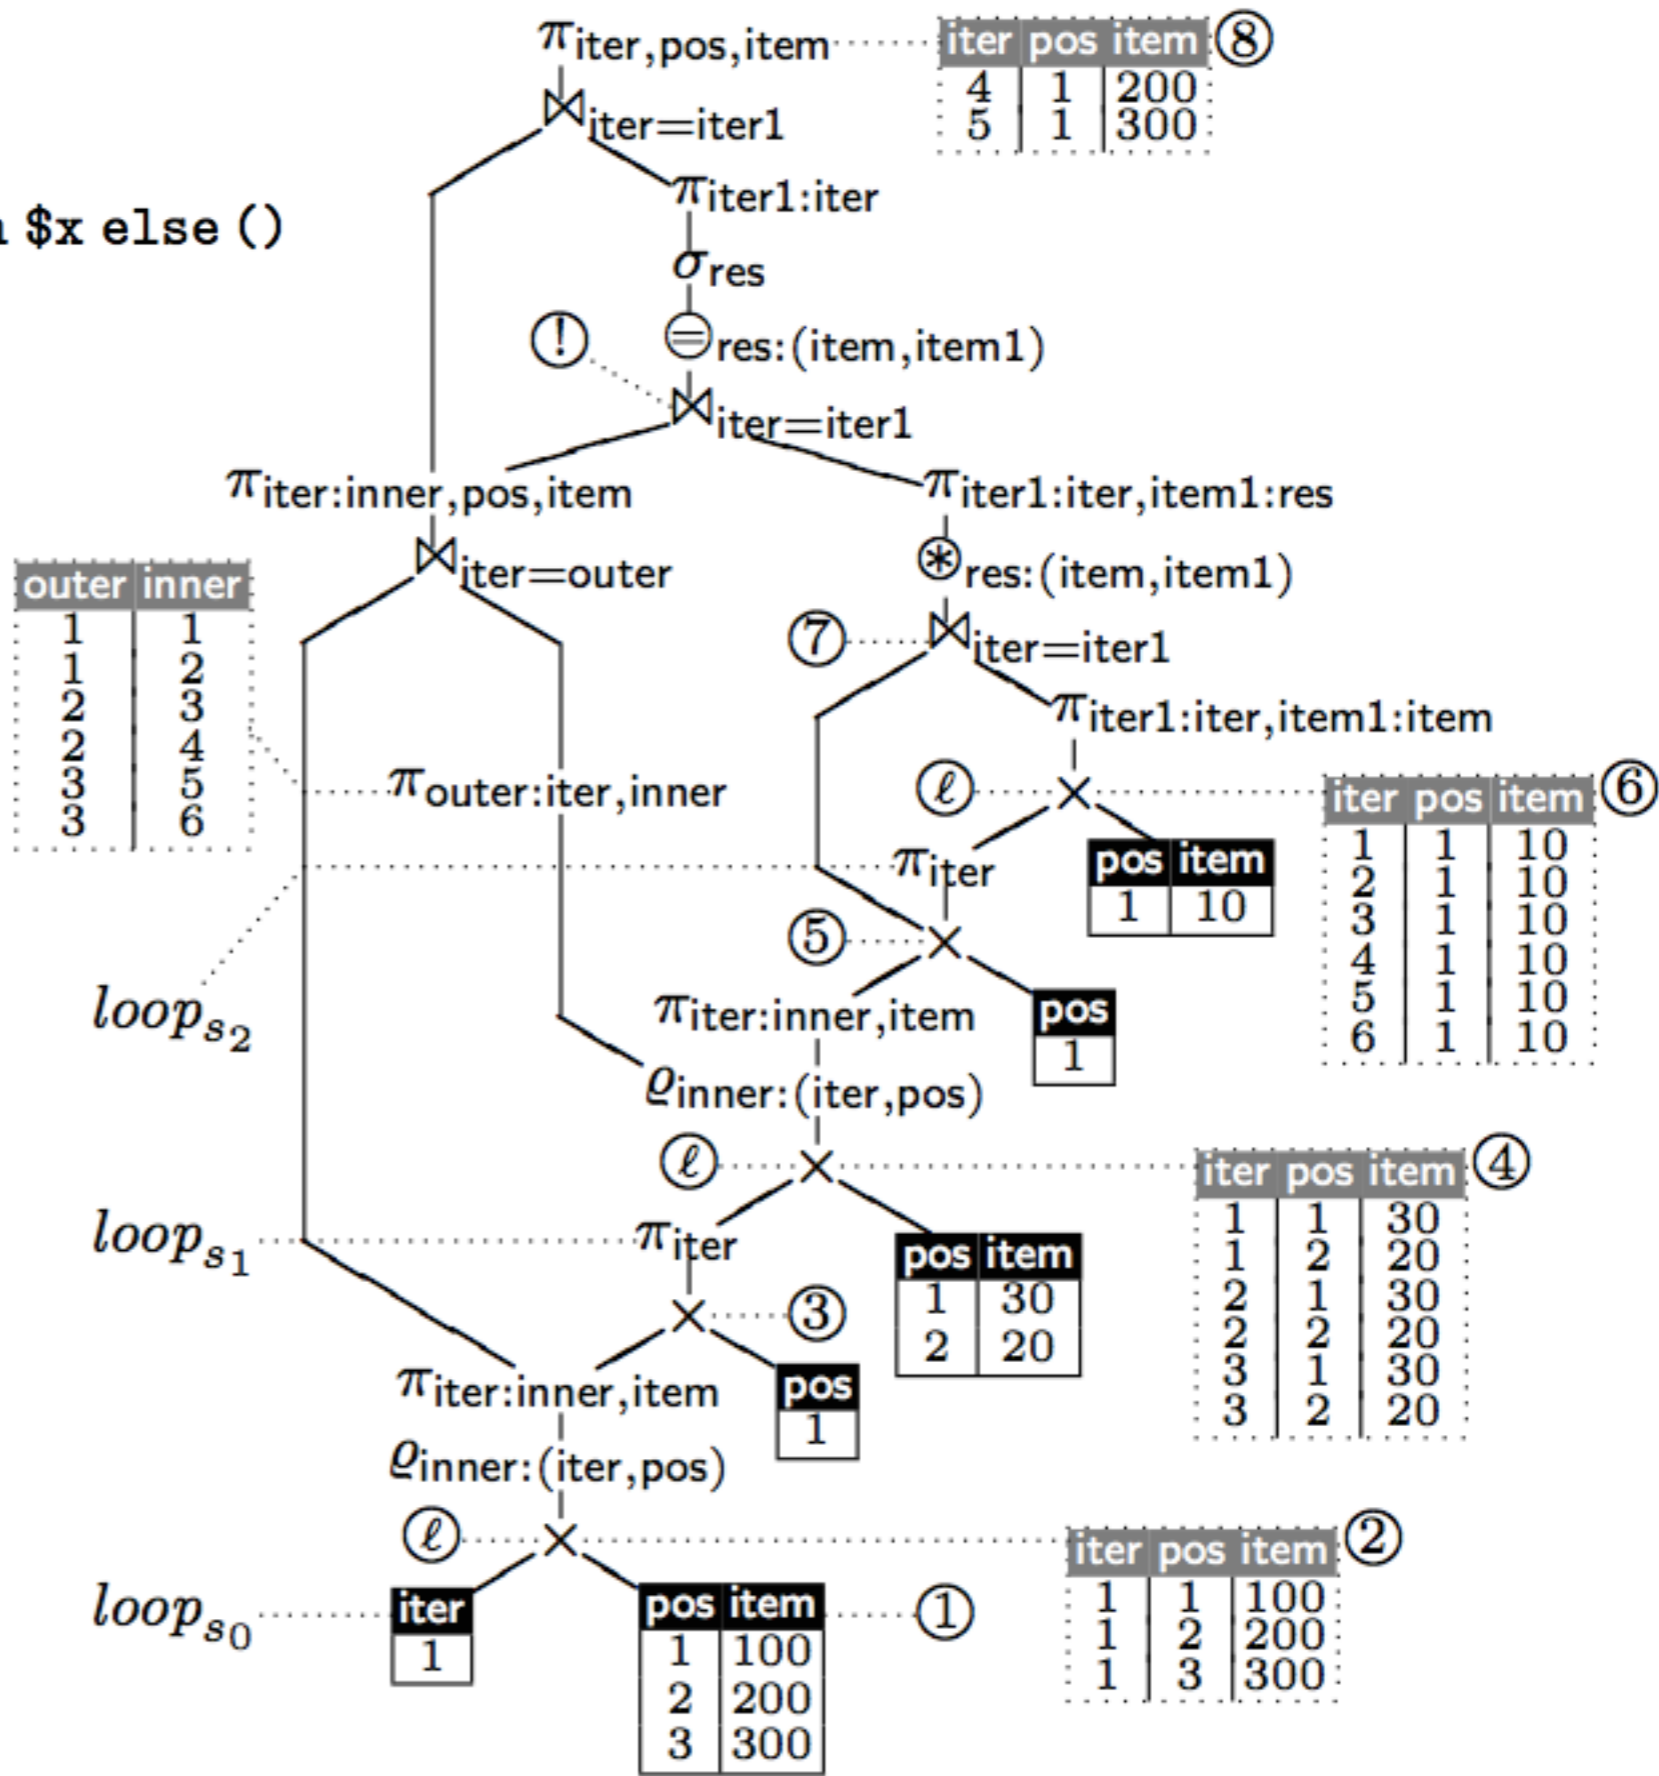
XIME-P 2005, June 16–17, Baltimore, Maryland.

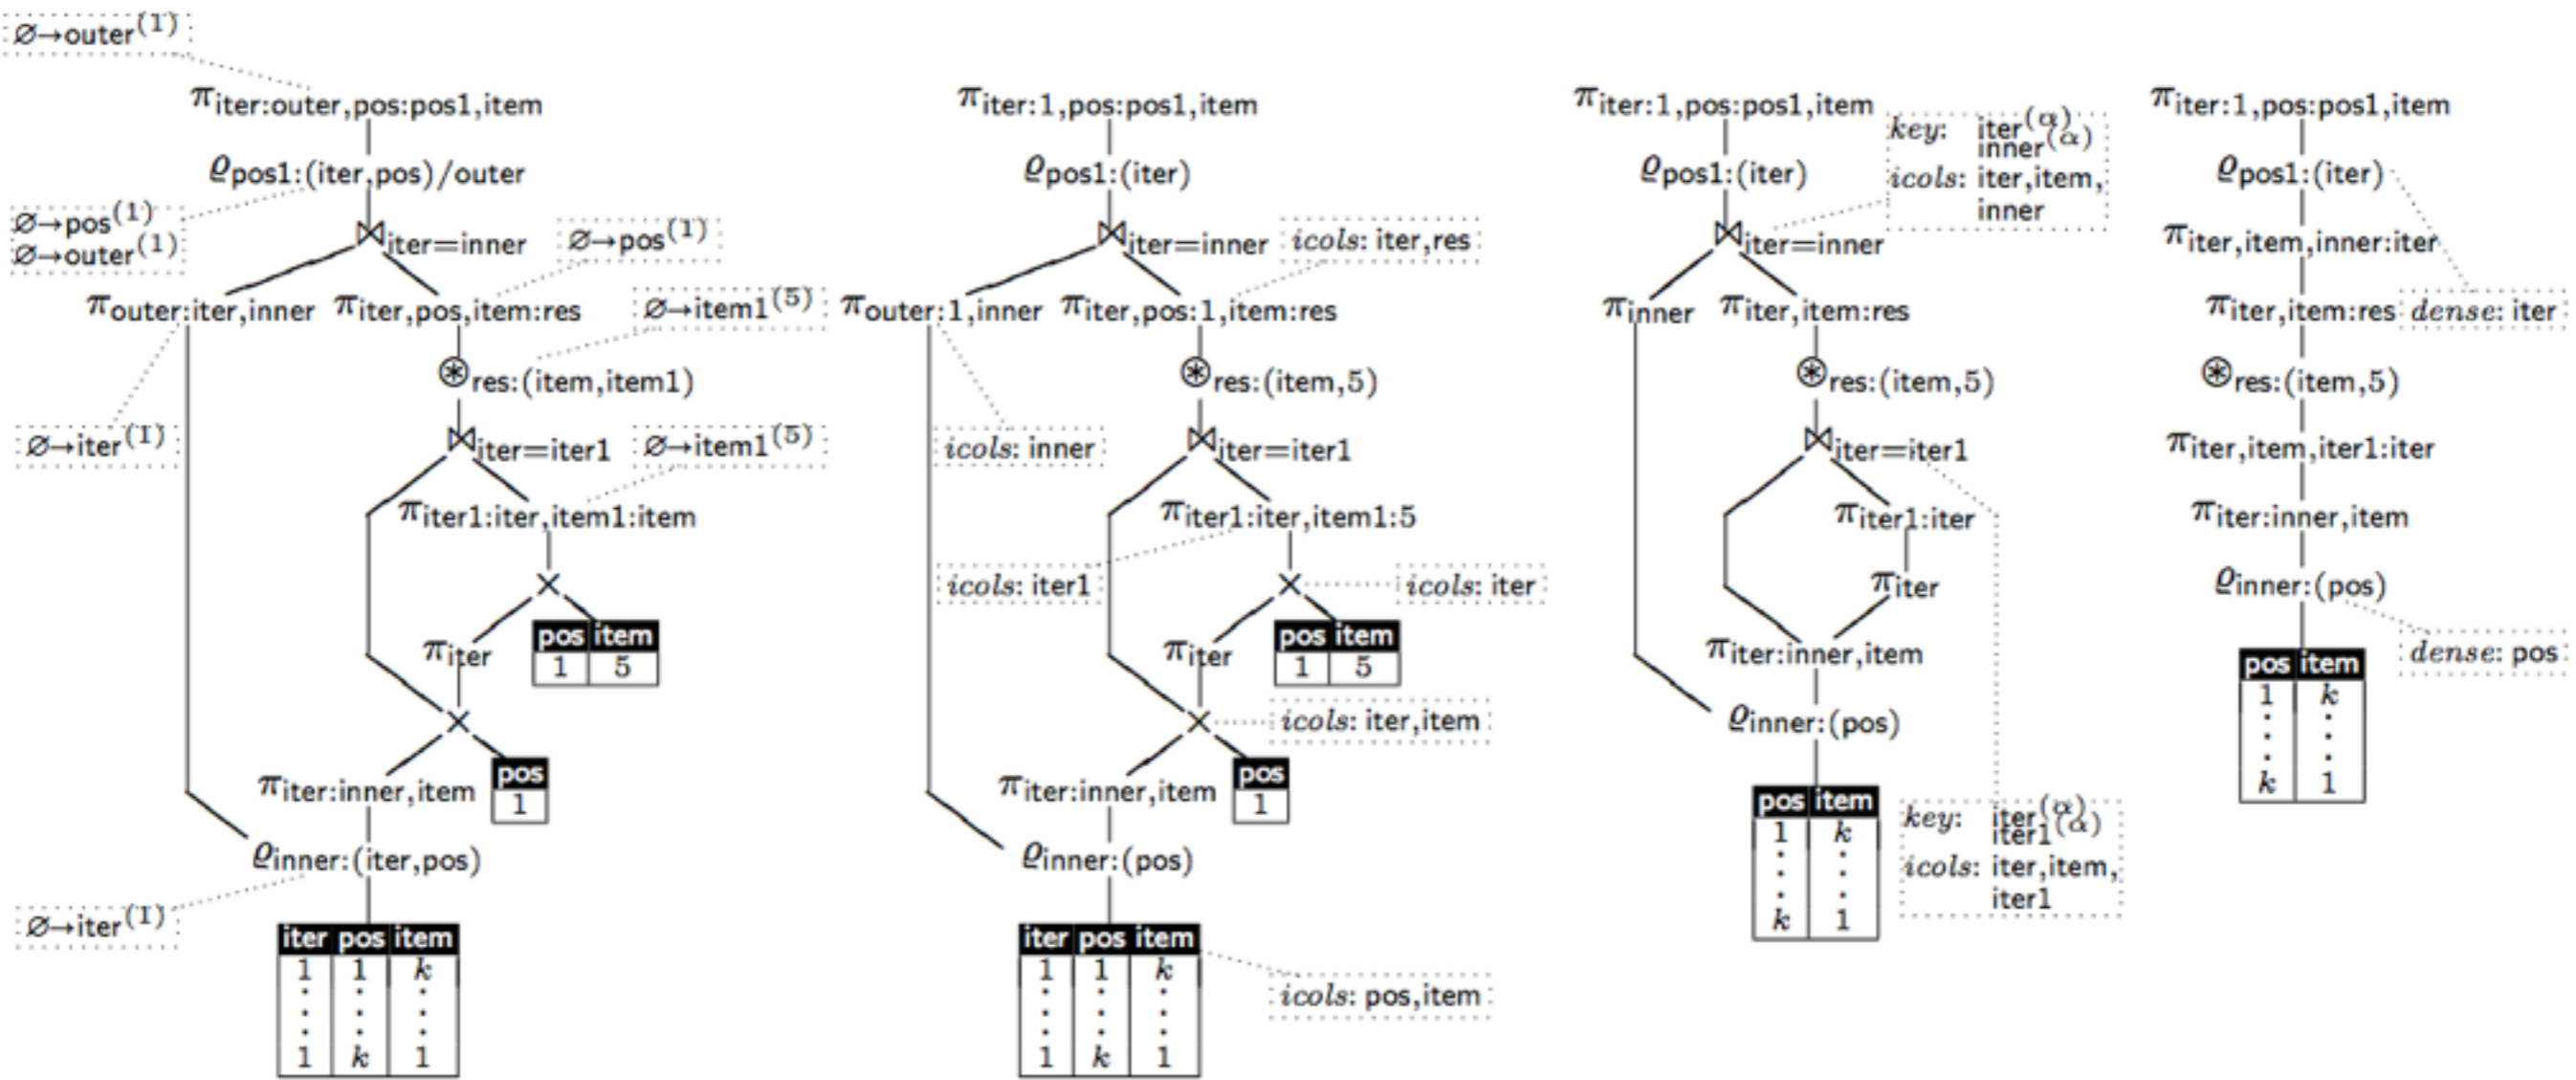
$$\begin{array}{l}
 s_0 \left\{ \begin{array}{l}
 \text{for } \$x \text{ in } \overbrace{(100, 200, 300)}^{e_1} \text{ return} \\
 s_1 \left\{ \begin{array}{l}
 \text{for } \$y \text{ in } \overbrace{(30, 20)}^{e_2} \text{ return} \\
 s_2 \left\{ \text{if } (\$x \text{ eq } \$y * \underbrace{10}_{e_3}) \text{ then } \$x \text{ else } () \right.
 \end{array} \right.
 \end{array} \right.
 \end{array}$$

s_0

```

    for $x in  $\overbrace{(100,200,300)}^{e_1}$  return
     $s_1$  {
      for $y in  $\overbrace{(30,20)}^{e_2}$  return
       $s_2$  { if ( $\overbrace{\$x \text{ eq } \$y * 10}^{e_3}$ ) then $x else () }
    }
  
```





9. REFERENCES

- [1] S. Boag, D. Chamberlin, M. Fernández, D. Florescu, J. Robie, and J. Simeon. XQuery 1.0: An XML Query Language. World Wide Web Consortium, Nov. 2003.
- [2] P. A. Boncz and M. L. Kersten. MIL Primitives for Querying a Fragmented World. *The VLDB Journal*, 8(2), 1999.
- [3] N. Bruno, N. Koudas, and D. Srivastava. Holistic Twig Joins: Optimal XML Pattern Matching. Madison, Wisconsin, USA, June 2002.
- [4] S. Chien, Z. Vagena, D. Zhang, V. Tsotras, and C. Zaniolo. Efficient Structural Joins on Indexed XML Documents. In *Proc. of the 28th Int'l Conference on Very Large Databases (VLDB)*, Hong Kong, China, Aug. 2002.
- [5] D. DeHaan, D. Toman, M. Consens, and M. Özsu. A Comprehensive XQuery to SQL Translation using Dynamic Interval Encoding. In *Proc. of the 22nd Int'l ACM SIGMOD Conference on Management of Data*, San Diego, California, USA, June 2003.
- [6] M. Fernández, J. Simeon, and P. Wadler. A Semi-monad for Semi-structured Data. In *Proc. of the 8th Int'l Conference on Database Theory (ICDT)*, London, UK, Jan. 2001.
- [7] D. Gluche, T. Grust, C. Mainberger, and M. Scholl. Incremental Updates for Materialized OQL Views. In *Proc. of the 5th Int'l Conference on Deductive and Object-Oriented Databases (DOOD)*, Montreux, Switzerland, Dec. 1997.
- [8] T. Grust. Accelerating XPath Location Steps. In *Proc. of the 21st Int'l ACM SIGMOD Conference on Management of Data*, Madison, Wisconsin, USA, June 2002.
- [9] T. Grust, S. Sakr, and J. Teubner. XQuery on SQL Hosts. In *Proc. of the 30th Int'l Conference on Very Large Data Bases (VLDB)*, Toronto, Canada, Aug. 2004.
- [10] 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, Sept. 2003.
- [11] J. Hidders and P. Michiels. Avoiding Unnecessary Ordering Operations in XPath. In *Proc. of the 9th Int'l Workshop on Database Programming Languages (DBPL)*, Potsdam, Germany, Sept. 2003.
- [12] R. Krishnamurthy, R. Kaushik, and J. Naughton. XML-to-SQL Query Translation Literature: The State of the Art and Open Problems. In *Proc. of the 1st Int'l XML Database Symposium (XSym)*, Berlin, Germany, Sept. 2003.
- [13] Q. Li and B. Moon. Indexing and Querying XML Data for Regular Path Expressions. In *Proc. of the 27th Int'l Conference on Very Large Databases (VLDB)*, Rome, Italy, Sept. 2001.
- [14] I. Manolescu, D. Florescu, and D. Kossmann. Answering XML Queries over Heterogeneous Data Sources. In *Proc. of the 27th Int'l Conference on Very Large Databases (VLDB)*, Rome, Italy, Sept. 2001.
- [15] J. Melton. *Advanced SQL:1999: Understanding Object-Relational and Other Advanced Features*. Morgan Kaufmann Publishers, Amsterdam, 2003.
- [16] A. Schmidt, F. Waas, M. Kersten, M. 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, Aug. 2002.
- [17] J. Shanmugasundaram, J. Kiernan, E. Shekita, C. Fan, and J. Funderburk. Querying XML Views of Relational Data. In *Proc. of the 27th Int'l Conference on Very Large Databases (VLDB)*, Rome, Italy, Sept. 2001.

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.

FERRY — Database-Supported Program Execution

Torsten Grust, Manuel Mayr, Jan Rittinger, and Tom Schreiber

WSI, Universität Tübingen
Tübingen, Germany

<firstname.lastname>@uni-tuebingen.de

```
1 let e = table Employees (id int, name string,  
2                          dept string, salary int)  
3     with keys ((id))  
4 in for x in e  
5     group by x.dept  
6     return (the (x.dept),  
7            take (2, for y in zip (x.name, x.salary)  
8            order by y.2 descending  
9            return y))
```

Copyright is held by the author/owner(s).

SIGMOD'09, June 29–July 2, 2009, Providence, Rhode Island, USA.

ACM 978-1-60558-551-2/09/06.

Haskell Boards the Ferry
















Database-Supported Program Execution for Haskell

George Giorgidze, Torsten Grust, Tom Schreiber, and Jeroen Weijers

Wilhelm-Schickard-Institut für Informatik,
Eberhard Karls Universität Tübingen

{george.giorgidze,torsten.grust,tom.schreiber,jeroen.weijers}@uni-tuebingen.de

2010

- [j8]    Torsten Grust: **Der Lehrstuhl für Datenbanksysteme am Wilhelm-Schickard-Institut der Universität Tübingen.** [Datenbank-Spektrum 10](#)(2): 105-106 (2010)
- [j7]    Torsten Grust, Jan Rittinger, Tom Schreiber: **Avalanche-Safe LINQ Compilation.** PVLDB 3(1): 162-172 (2010)
- [j6]    Tom Schreiber, Simone Bonetti, Torsten Grust, Manuel Mayr, Jan Rittinger: **Thirteen New Players in the Team: A Ferry-based LINQ to SQL Provider.** PVLDB 3(2): 1549-1552 (2010)
- [c30]    Torsten Grust, Manuel Mayr, Jan Rittinger: **Let SQL drive the XQuery workhorse (XQuery join graph isolation).** EDBT 2010: 147-158
- [c29]    George Giorgidze, Torsten Grust, Tom Schreiber, Jeroen Weijers: **Haskell Boards the Ferry - Database-Supported Program Execution for Haskell.** IFL 2010: 1-18

A study of the exact relationship between DPH and DSH still lies ahead. We conjecture that DSH's loop-lifting compilation strategy does have an equivalent formulation in terms of vectorisation or Blelloch's flattening transformation [4].

```

fst^      :: [:(a, b):] → [:(a):]
snd^     :: [:(a, b):] → [:(b):]
bpermuteP :: [:(a):] → [:(Int):] → [:(a):]
let sv = [:(1, 0.1), (3, 1.0), (4, 0.0):]
      v = [:(10, 20, 30, 40, 50):]
in
      sumP (snd^ sv *^ bpermuteP v (fst^ sv))

```

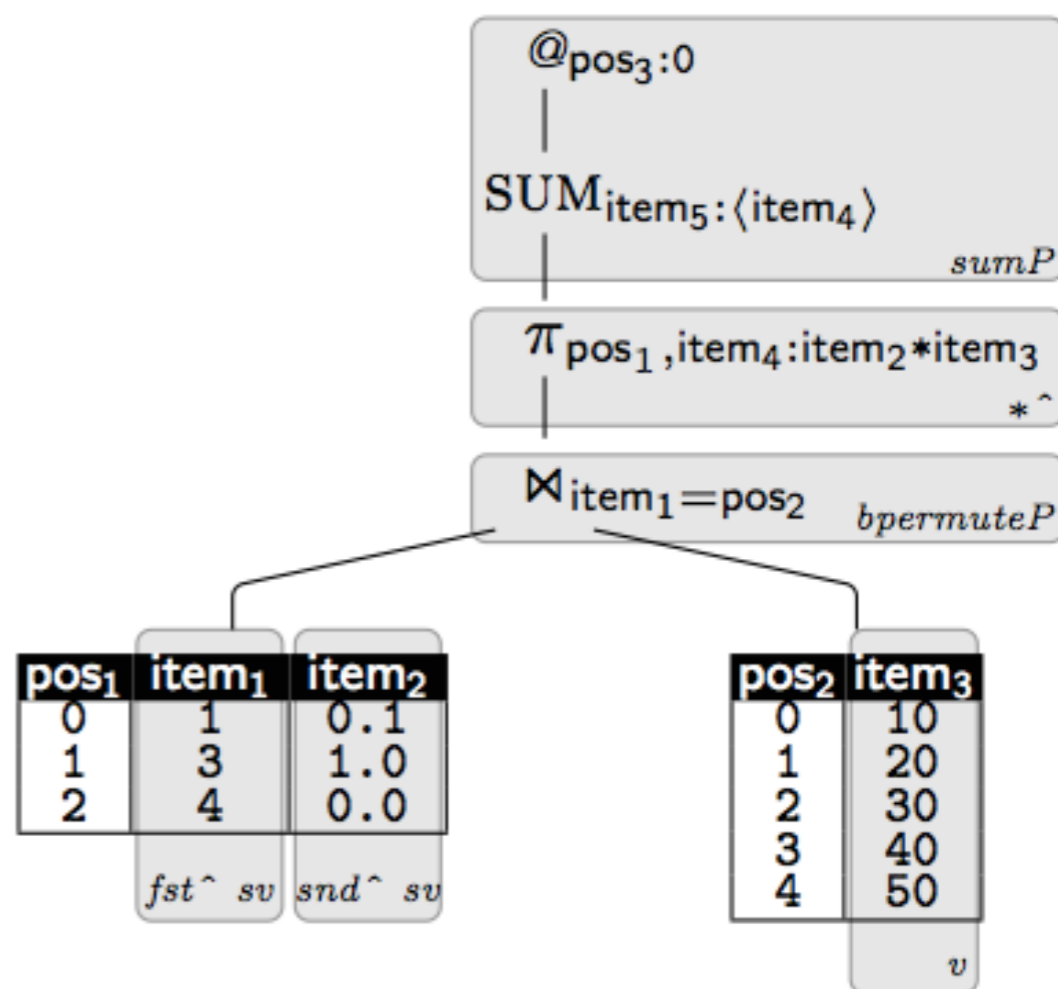


Fig. 6. Intermediate code generated for the sparse vector multiplication example of Fig. 5: DPH (left) *vs.* DSH (right).