

Dependable Cardinality Forecasts for XQuery

Jens Teubner, ETH (formerly IBM Research)
Torsten Grust, U Tübingen (formerly TUM)
Sebastian Maneth, UNSW and NICTA
Sherif Sakr, UNSW and NICTA



Cardinality Estimation for XQuery

The feature-richness and semantics of the language make cardinality estimation for XQuery notoriously hard.

```
for $d in doc("forecast.xml")/descendant::day
let $day := $d/@t
let $ppcp := data($d/descendant::ppcp)
return
  if ($ppcp > 50)
    then ("rain likely on", $day,
         "chance of precipitation:", $ppcp)
    else ("no rain on", $day)
```

- ▶ for iteration
 - ▶ conditionals (if-then-else)
 - ▶ existential quantification
 - ▶ sequence construction
 - ▶ ...
- (XPath is **not** a focus of this work.)

Cardinality Estimation for XQuery

The feature-richness and semantics of the language make cardinality estimation for XQuery notoriously hard.

```

for $d in doc("forecast.xml")/descendant::day card1 = ?
let $day := $d/@t
let $ppcp := data($d/descendant::ppcp)
return
  if ($ppcp > 50) card2 = ?
  then ("rain likely on", $day, card3 = ?
        "chance of precipitation:", $ppcp)
  else ("no rain on", $day) card4 = ?

```

- ▶ for iteration
- ▶ conditionals (if-then-else)
- ▶ existential quantification
- ▶ sequence construction
- ▶ ...
- (XPath is **not** a focus of this work.)

→ **Goal:** Compute subexpression-level cardinalities $card_j$.

Idea: Perform cardinality estimation on relational plan equivalents for XQuery.

relational cardinality estimation (System R)
+ existing work on XPath estimation
+ histograms for value predicates
—
= cardinality estimation for XQuery

- ▶ Build on Pathfinder's XQuery-to-relational algebra compiler.¹
 - tuple count \equiv XQuery item count
- ▶ Cardinality information for **each** subexpression.

¹<http://www.pathfinder-xquery.org/>

Relational XQuery Cardinality Estimation

Apply System R-style estimation to relational XQuery plans, *e.g.*,

Disjoint union:

$$|q_1 \cup q_2| = |q_1| + |q_2|$$

Cartesian product:

$$|q_1 \times q_2| = |q_1| \cdot |q_2|$$

Equi-join:

$$|q_1 \bowtie_{a=b} q_2| = \begin{cases} \frac{|q_1| \cdot |q_2|}{\max\{|a|_{\text{idx}}, |b|_{\text{idx}}\}} & \text{if there are indexes on} \\ & \text{both join columns,} \\ \frac{|q_1| \cdot |q_2|}{|a|_{\text{idx}}} & \text{if there is only an index} \\ & \text{on column } a, \\ |q_1| \cdot |q_2| \cdot 1/10 & \text{otherwise} \end{cases}$$

$|c|_{\text{idx}}$: Number of unique values in index on column c .

Relational XQuery Cardinality Estimation

Apply System R-style estimation to relational XQuery plans, *e.g.*,

Disjoint union:

$$|q_1 \cup q_2| = |q_1| + |q_2|$$

Cartesian product:

$$|q_1 \times q_2| = |q_1| \cdot |q_2|$$

Equi-join:

$$|q_1 \bowtie_{a=b} q_2| = \begin{cases} \frac{|q_1| \cdot |q_2|}{\max\{|a|_{\text{idx}}, |b|_{\text{idx}}\}} & \text{if there are indexes on both join columns,} \\ \frac{|q_1| \cdot |q_2|}{|a|_{\text{idx}}} & \text{if there is only an index on column } a, \\ |q_1| \cdot |q_2| \cdot 1/10 & \text{otherwise} \end{cases}$$

- Our joins typically operate over **computed** relations.

$|c|_{\text{idx}}$: Number of unique values in index on column c .

Abstract Domain Identifiers

A simple form of **data flow analysis** provides the information needed.

- ▶ Introduce **abstract domain identifiers** α, β, \dots as placeholders for the **active runtime domain** for each column c .
(Read c^α as “column c contains values from domain α .”)
- ▶ Estimate the **size** $\|\alpha\|$ of each domain α , e.g.,²

$$\text{dom}(\rho_{a:\langle b_1, \dots, b_n \rangle}(q)) \supseteq \text{dom}(q) \cup \left\{ a^\alpha \wedge \|\alpha\| = |q| \right\} .$$

- ▶ Identify **inclusion relationships** $\alpha \sqsubseteq \beta$ between domains, e.g.,

$$a^\alpha \in \text{dom}(q) \wedge a^\beta \in \text{dom}(\sigma\dots(q)) \Rightarrow \beta \sqsubseteq \alpha .$$

²Operator $\rho_{a:\langle b_1, \dots, b_n \rangle}$ introduces a new key column (holding row numbers).

Abstract Domain Identifiers

Use abstract domain information for cardinality estimation.

E.g., “foreign key” join:

$$\frac{a^\alpha \in \text{dom}(q_1) \quad b^\beta \in \text{dom}(q_2) \quad \alpha \sqsubseteq \beta}{|q_1 \bowtie_{a=b} q_2| = \frac{|q_1| \cdot |q_2|}{\|\beta\|}}$$


- ▶ Domain inclusion guarantees that each tuple in q_1 finds (at least one) join partner in q_2 .

Other examples:

- ▶ $|q_1 \setminus q_2| = |q_1| - |q_2|$ if q_2 is a subset of q_1 .
- ▶ $|q_1 \setminus q_2| = 0$ if q_1 is a subset of q_2 .
- ▶ $|q_1 \setminus q_2| = |q_1|$ if q_1 and q_2 are disjoint.

Interfacing with XPath—Projection Paths


Track XPath navigation by means of **projection paths**³



- ▶ $b \Rightarrow^p \in \text{path}(q_1) \Rightarrow c \Rightarrow^{p/\text{child}::*} \in \text{path}(q_2)$
- ▶ Step operator \sqsubset makes XPath navigation explicit in relational plans (compiles to join on SQL back-ends).

³A. Marian and J. Siméon. Projecting XML Documents. *VLDB 2003*.

Interfacing with XPath—Cardinality Inference



Cardinality:

$$|q_2| = |q_1| \cdot \frac{\text{fn:count}(p/\text{child}::*)}{\text{fn:count}(p)} = |q_1| \cdot \underbrace{\text{Pr}_{\text{child}::*}(p)}_{\text{fanout (here: } 4/3)}$$


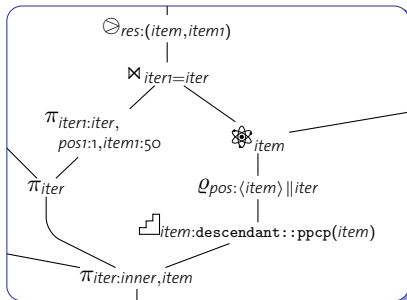
Domain Sizes:

$$\|\alpha'\| = \|\alpha\| \cdot \frac{\text{fn:count}(p[\text{child}::*])}{\text{fn:count}(p)} = \|\alpha\| \cdot \underbrace{\text{Pr}_{[\text{child}::*]}(p)}_{\text{selectivity (here: } 2/3)}$$

Any XPath estimator that provides $\text{Pr}_{p_2}(p_1)$ and $\text{Pr}_{[p_2]}(p_1)$ will do.


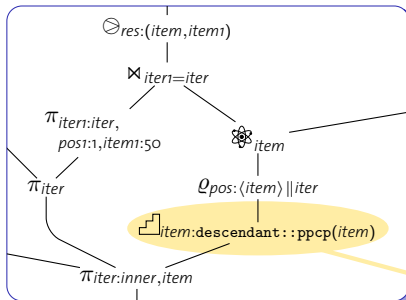
- Our prototype uses a simple Data Guide-based implementation.

Back to our Example Plan



\otimes_a : Retrieve **typed values** for node identifiers in column a (atomization).

Back to our Example Plan




Projection path:

$item \Rightarrow \dots / descendant :: ppcp \in path(\sqcup \dots (q))$

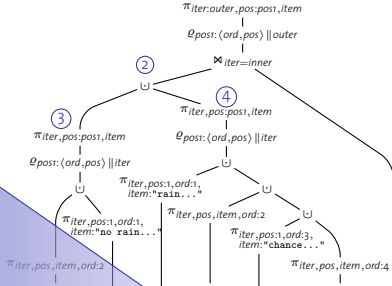
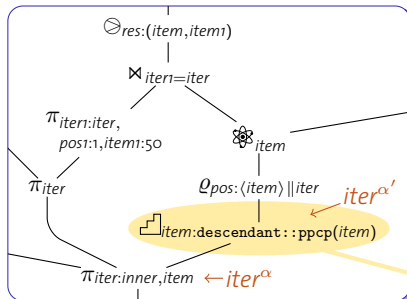
Cardinality (uses fanout):

$|\sqcup item:ax::nt(item)(q)| = |q| \cdot Pr_{ax::nt}(\dots)$

\otimes_a : Retrieve **typed values** for node identifiers in column a (atomization).



Back to our Example Plan



Projection path:

$item \Rightarrow \dots / descendant :: ppcp \in path(\sqcup \dots(q))$

Cardinality (uses fanout):

$|\sqcup_{item:ax::nt}(q)| = |q| \cdot Pr_{ax::nt}(\dots)$


Domain inclusion:

$iter^{\alpha'} \in dom(\sqcup \dots(q)); \alpha' \sqsubseteq \alpha$


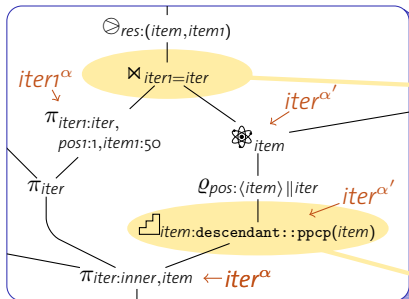
Domain size (uses step selectivity):

$\| \alpha' \| = \| \alpha \| \cdot Pr_{[ax::nt]}(\dots)$

\otimes_a : Retrieve **typed values** for node identifiers in column a (atomization).



Back to our Example Plan



“Foreign key” join:

$$|q_1 \bowtie_{iter=iter} q_2| = \frac{|q_1| \cdot |q_2|}{\|\alpha\|} \quad (\text{since } \alpha' \sqsubseteq \alpha)$$

Projection path:

$$item \Rightarrow \dots / descendant :: ppcp \in path(\lceil \dots (q))$$

Cardinality (uses fanout):

$$|\lceil_{item:ax::nt}(q)| = |q| \cdot Pr_{ax::nt}(\dots)$$


Domain inclusion:

$$iter^{\alpha'} \in dom(\lceil \dots (q)); \alpha' \sqsubseteq \alpha$$

Domain size (uses step selectivity):

$$\|\alpha'\| = \|\alpha\| \cdot Pr_{[ax::nt]}(\dots)$$

\otimes_a : Retrieve **typed values** for node identifiers in column a (atomization).

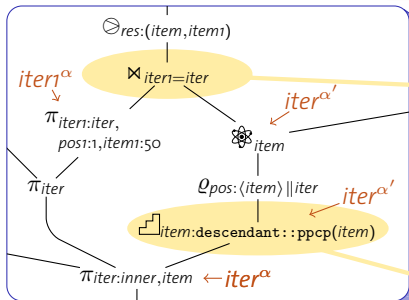


Back to our Example Plan

card₂: 4104

card₃: 3540

card₄: 564



“Foreign key” join:

$$|q_1 \bowtie_{iter=iter} q_2| = \frac{|q_1| \cdot |q_2|}{\|\alpha\|} \quad (\text{since } \alpha' \sqsubseteq \alpha)$$

Projection path:

$$item \Rightarrow \dots / descendant :: ppcp \in path(\sqcup \dots (q))$$

Cardinality (uses fanout):

$$|\sqcup_{item:ax::nt(item)}(q)| = |q| \cdot Pr_{ax::nt}(\dots)$$


Domain inclusion:

$$iter^{\alpha'} \in dom(\sqcup \dots (q)); \alpha' \sqsubseteq \alpha$$

Domain size (uses step selectivity):

$$\|\alpha'\| = \|\alpha\| \cdot Pr_{[ax::nt]}(\dots)$$

\otimes_a : Retrieve **typed values** for node identifiers in column a (atomization).



Forecasting New Zealand's Weather

The obtained cardinalities can be mapped back to predict item counts for corresponding XQuery expressions:⁴

```
for $d in doc("forecast.xml")/descendant::day card1 = 990/990
let $day := $d/@t
let $ppcp := data($d/descendant::ppcp)
return
  if ($ppcp > 50) card2 = 4104/3402
  then ("rain likely on", $day, card3 = 3540/2370
        "chance of precipitation:", $ppcp)
  else ("no rain on", $day) card4 = 564/1032
```

- ▶ Value statistics based on histograms (↗ paper)
- ▶ Inaccuracy is mainly due to correlations in the data.
 - ▶ Rain in the morning likely means rain in the afternoon, too.


⁴estimated/observed, based on data taken for New Zealand two weeks ago.

More Realistic Queries: W3C XQuery Use Cases

- ▶ Prototype implementation based on Pathfinder
- ▶ For each subexpression:

$$\frac{\text{estimated cardinality}}{\text{observed cardinality}}$$
- ▶ Diameter indicates data point “stacking”
- ▶ Plan root: filled circle ●
- ▶ Applicable to “real” queries
- ▶ Recovery from intermediate mis-estimations
 - ▶ e.g., existential semantics

XMP Q4
XMP Q5
XMP Q6
XMP Q8
XMP Q9
XMP Q10
XMP Q11
SGML Q3
SGML Q4
SGML Q8a
SGML Q8b
R Q2
R Q3
R Q10
R Q11
R Q13
R Q14
R Q15
R Q17



Wrap-Up

- ▶ Cardinality estimation framework for XQuery
- ▶ **Subexpression-level** estimates for **arbitrary** XQuery expressions
- ▶ Based on Pathfinder's XQuery-to-relational algebra compiler

relational cardinality estimation (System R)
+ existing work on XPath estimation
+ histograms for value predicates

= cardinality estimation for XQuery

- ▶ High-quality estimates for realistic XQuery workloads
 - ▶ **Robust** with respect to intermediate errors
- ▶ **Pluggable** and **extensible**
 - ▶ *e.g.*, XPath estimation subsystem, positional predicates