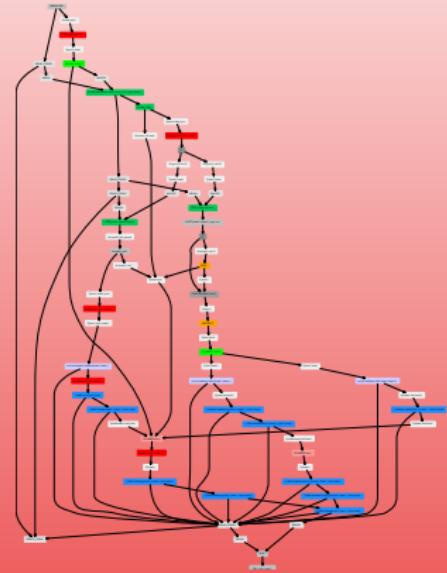


Pathfinder: A Relational Query Optimizer Explores XQuery Terrain

Torsten Grust · Jan Rittinger · Jens Teubner · TU München



Pathfinder consumes XQuery with arbitrary expression nestings ...

literals	42, "foo", (), ...
arithmetics	$e_1 + e_2, e_1 - e_2, \dots$
built-in functions	<code>fn:sum(e), fn:count(e), fn:doc(uri), ...</code>
variable bindings	<code>let \$v := e₁ return e₂</code>
iteration	<code>for \$v at \$p in e₁ return e₂</code>
conditionals	<code>if p then e₁ else e₂</code>
sequence construction	e_1, e_2
user-defd. functions	$f(e_1, e_2, \dots, e_n)$
element construction	<code>element e₁ { e₂ }</code>
XPath steps	$e/\alpha::\nu$ (full axis feature)
:	:

Pathfinder consumes XQuery with arbitrary expression nestings ...

literals
arithmetics
built-in functions
variable bindings
iteration
conditionals
sequence construction
user-defd. functions
element construction
XPath steps
⋮

42, "foo", (), ...
 $e_1 + e_2, e_1 - e_2, \dots$
`fn:sum(e), fn:count(e), fn:doc(uri), ...`
`let $v := e1 return e2`
`for $v at $p in e1 return e2`
`if $p then e1 else e2`
 e_1, e_2
`(e1, e2, ..., en)`
`element e1 { e2 }`
`e/α::ν (full axis feature)`
⋮

FULL XQUERY SUPPORT

... and compiles them into plans of a standard relational algebra.

π	column projection, renaming
σ	row selection
\bowtie	equi-join
\times	Cartesian product
\sqcup, \setminus	disjoint union, difference
δ	duplicate elimination
ϱ	row numbering
	staircase join ¹
ε, τ	element/text node construction ¹
\circledast	arithmetic/comparison/Boolean operator *

- Operates on node (not tree!) level and 1NF relations.

¹Syntactic sugar; expressible by remaining operators.

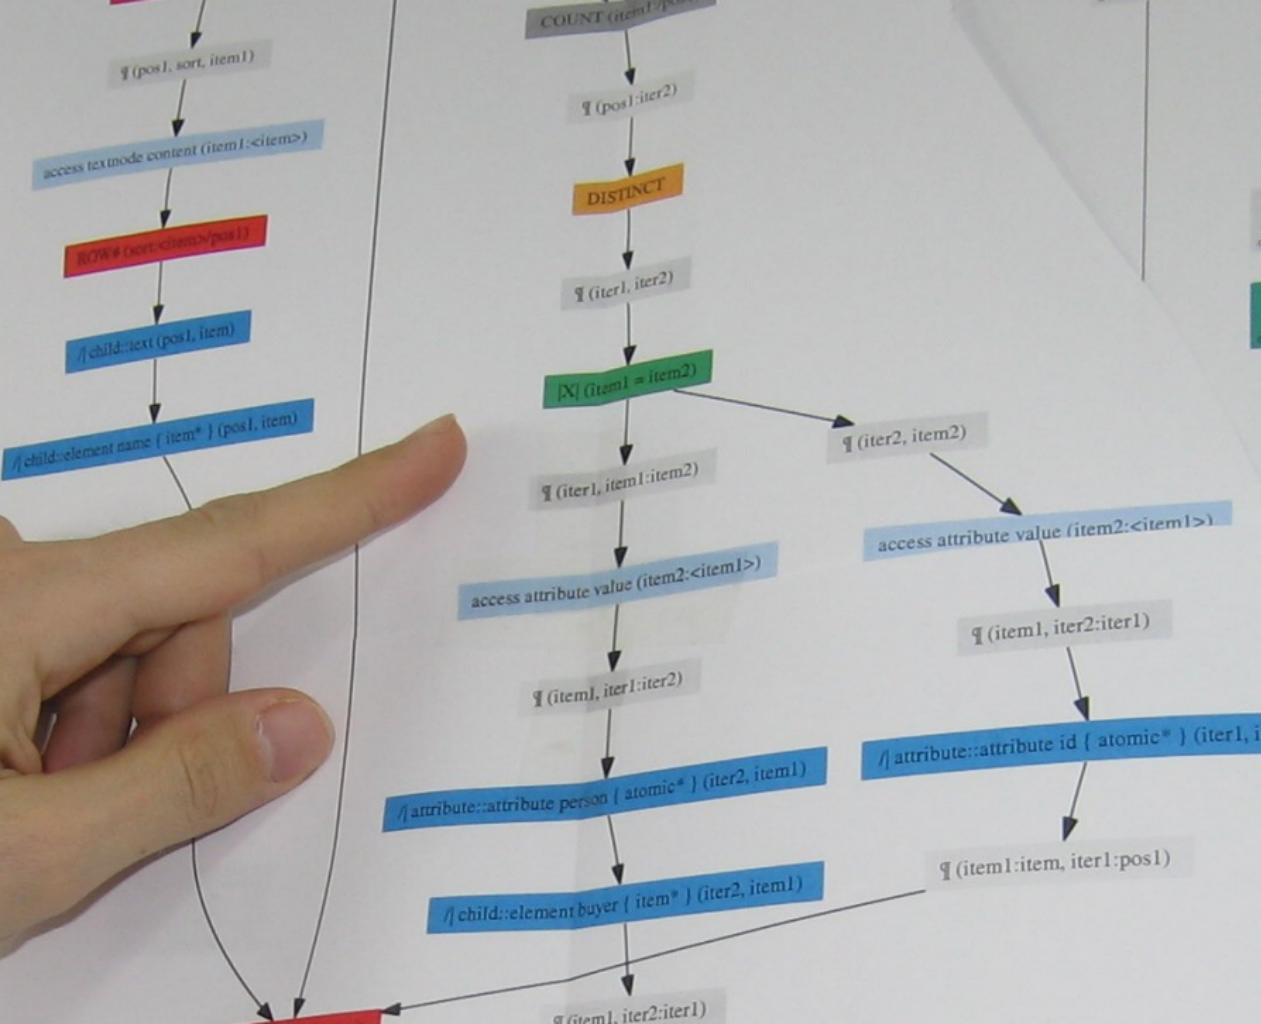
... and compiles them into plans of a standard relational algebra.

π	column projection, renaming
σ	row selection
\bowtie	equi-join
\times	Cartesian product
\sqcup, \setminus	disjoint union, difference
δ	duplicate elimination
ϱ	row numbering
	staircase join
ε, τ	element/text node construction ¹
$*$	arithmetic/comparison/Boolean operator *

EFFICIENTLY
IMPLEMENTABLE
ON SQL 99 HOSTS

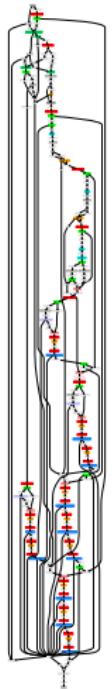
- ▶ Operates on node (not tree!) level and 1NF relations.

¹Syntactic sugar; expressible by remaining operators.





To combat the resulting plan sizes, Pathfinder uses



Resulting query runtime (MonetDB/XQuery):

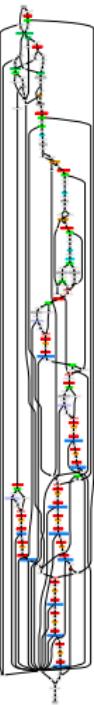
0

112 sec

To combat the resulting plan sizes, Pathfinder uses

- 1 constant propagation,

Resulting query runtime (MonetDB/XQuery):

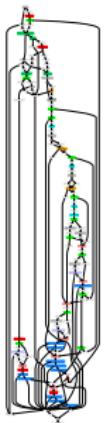


1 0

103 sec 112 sec

To combat the resulting plan sizes, Pathfinder uses

- 1 constant propagation,
- 2 projection pushdown,

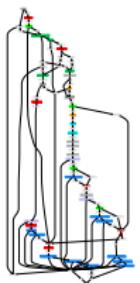


Resulting query runtime (MonetDB/XQuery):



To combat the resulting plan sizes, Pathfinder uses

- 1 constant propagation,
- 2 projection pushdown,
- 3 functional dependency and data flow analyses, and

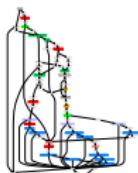


Resulting query runtime (MonetDB/XQuery):



To combat the resulting plan sizes, Pathfinder uses

- 1 constant propagation,
- 2 projection pushdown,
- 3 functional dependency and data flow analyses, and
- 4 algebraic join detection.



Resulting query runtime (MonetDB/XQuery):



Get your copy of Pathfinder today:

<http://www.pathfinder-xquery.org/>

Pathfinder is developed at TU München by

- **Torsten Grust, Jan Rittinger, and Jens Teubner.**

Thanks to our long-term partners

- Peter Boncz, CWI Amsterdam
- Stefan Manegold, CWI Amsterdam
- Jan Flokstra, U Twente
- Sjoerd Mullender, CWI Amsterdam
- Maurice van Keulen, U Twente
- Ying Zhang, CWI Amsterdam
- Martin Kersten, CWI Amsterdam

Pathfinder is supported by the German Research Foundation **DFG**.