

Unnesting Arbitrary Queries

Thomas Neumann and Alfons Kemper

Technische Universität München

March 5, 2015

Often queries are simpler to formulate using subqueries

```
Q1: select s.name,e.course
      from  students s,exams e
      where s.id=e.sid and
            e.grade=(select min(e2.grade)
                      from exams e2
                      where s.id=e2.sid)
```

- here, subquery depends on outer query (correlated)
- nested loop evaluation, $O(n^2)$
- easy to formulate, very inefficient to execute!

Same query without correlated subquery:

```
Q1': select s.name,e.course
       from students s,exams e,
           (select e2.sid as id, min(e2.grade) as best
            from exams e2
            group by e2.sid) m
       where s.id=e.sid and m.id=s.id and
           e.grade=m.best
```

- much more efficient to execute, no longer $O(n^2)$
- but not as intuitive as the original query
- a database should *unnest* (i.e., de-correlate) automatically

Motivation (3)

Typically, DBMSs detect and unnest some simple cases. But correlations can be complex:

Q2:

```
select s.name, e.course
from   students s, exams e
where  s.id=e.sid and
       (s.major = 'CS' or s.major = 'Games Eng') and
       e.grade>=(select avg(e2.grade)+1
                  from exams e2
                  where s.id=e2.sid or
                       (e2.curriculum=s.major and
                        s.year>e2.date))
```

- “difficult” (non-equality, disjunction, etc.)
- we are not aware of any system that could unnest that
- but $O(n^2)$ is a deal breaker, a DBMS must avoid that if possible

SQL promised declarative queries

- the user writes what he wants, not what the system should do
- the DBMS finds a good (the best?) evaluation strategy
- failing to unnest queries often leads to catastrophic runtime

We want an generic approach that can handle arbitrary queries

- works on the algebra, on on the SQL representation
- can handle all relational operators

We need some extra functionality

$$\chi_{a:f}(e) := \{x \circ (a : f(x)) \mid x \in e\}$$

$$T_1 \bowtie_p T_2 := \sigma_p(T_1 \times T_2)$$

$$T_1 \Join_p T_2 := \{t_1 \circ t_2 \mid t_1 \in T_1 \wedge t_2 \in T_2(t_1) \wedge p(t_1 \circ t_2)\}$$

$$\Gamma_{A;a:f}(e) := \{x \circ (a : f(y)) \mid x \in \Pi_A(e) \wedge y = \{z \mid z \in e \wedge \forall a \in A : x.a = z.a\}\}$$

Additional notation:

$$\mathcal{A}(T) := \text{the attributes produced by } T$$

$$\mathcal{F}(T) := \text{the free variables of } T$$

Canonical translation turns correlated subqueries into

$$(outer\ query) \bowtie_p (subquery).$$

- \bowtie is a dependent join (evaluates right hand side for every tuple)
- nested loop evaluation, very expensive

The goal of unnesting is to eliminate all dependent joins.

Some cases are simple

```
select ...
from   lineitem l1 ...
where  exists (select *
               from lineitem l2
               where l2.l_orderkey = l1.l_orderkey)
...

```

This results in an algebra expression of the form

$$l_1 \bowtie (\sigma_{l_1.okey=l_2.okey}(l_2))$$

We can unnest by pulling the predicate up, eliminating the dependency.

$$l_1 \bowtie_{l_1.okey=l_2.okey} (l_2)$$

- pull predicates up to eliminate correlations


General idea: Evaluate subquery for all possible bindings simultaneously.

$$T_1 \bowtie_p T_2 \equiv T_1 \bowtie_{p \wedge T_1 = \mathcal{A}(D)} D (D \bowtie T_2)$$

where $D := \Pi_{\mathcal{F}(T_2) \cap \mathcal{A}(T_1)}(T_1)$.

- D provides all possible bindings of free variables
- $|D| \leq |T_1|$
- D is a set (i.e., duplicate free)
- D being a set allow for equivalence that do not hold in general
- allows us to move D until subquery no longer dependent

General Unnesting (2)




Using D might already improve runtime sometimes, but in general is only the first step for full unnesting.

A dependent join with a set D can be manipulated much more easily. We push D down until the join is no longer dependent:


$$D \bowtie T \equiv D \bowtie T \quad \text{if} \quad \mathcal{F}(T) \cap \mathcal{A}(D) = \emptyset.$$

Push down rules vary between operators:

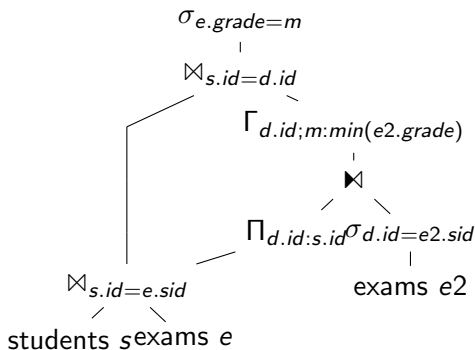
$$\begin{aligned}
 D \bowtie \sigma_p(T_2) &\equiv \sigma_p(D \bowtie T_2) \\
 D \bowtie (T_1 \bowtie_p T_2) &\equiv \begin{cases} (D \bowtie T_1) \bowtie_p T_2 & : \mathcal{F}(T_2) \cap \mathcal{A}(D) = \emptyset \\ T_1 \bowtie_p (D \bowtie T_2) & : \mathcal{F}(T_1) \cap \mathcal{A}(D) = \emptyset \\ (D \bowtie T_1) \bowtie_{p \wedge \text{natural } D} (D \bowtie T_2) & : \text{otherwise.} \end{cases} \\
 D \bowtie (T_1 \bowtie_{p \wedge \text{natural } D} T_2) &\equiv (D \bowtie T_1) \bowtie_{p \wedge \text{natural } D} (D \bowtie T_2) \\
 D \bowtie (\Gamma_{A;a:f}(T)) &\equiv \Gamma_{A \cup \mathcal{A}(D);a:f}(D \bowtie T) \\
 \dots &\quad \text{(see the paper)}
 \end{aligned}$$



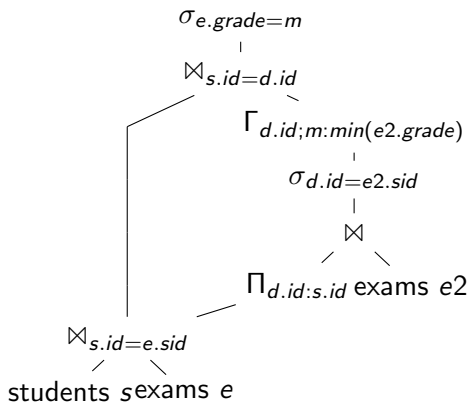
Original Query 1



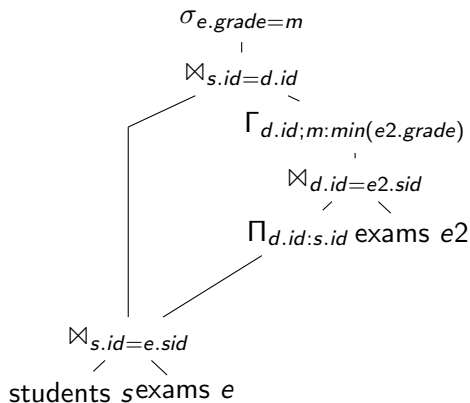
Query 1, Transformation Step 1



Query 1, Transformation Step 2



Query 1, Transformation Step 4




Query 1, Transformation Step 5 (pushing selections back down)

Instead of joining with D , we can often *infer* the attributes from D

$$D \bowtie T \subseteq \chi_{\mathcal{A}(D):B}(T) \text{ if } \exists B \subseteq \mathcal{A}(T) : \mathcal{A}(D) \equiv_C B.$$

- “perfect” unnesting, totally independent query parts afterwards
- but: this computes a *superset* of the join with D
- does not matter for correctness (final join will eliminate non- D values), but for performance
- we avoid computing D , but we potential lose pruning power
- a good idea if the join is unselective, otherwise keep D
- cost-base decision



Query 1, Optional Transformation Step 6 (decoupling both sides)

- unnesting transforms an $O(n^2)$ into an (ideally) $O(n)$ operation
- arbitrary gains possible

Toy database, 1,000 students, 10,000 exams (i7-3930K)

	Q1	Q2
HyPer	< 1ms	42ms
HyPer without unnesting	51ms	408ms
PostgreSQL 9.1	1,300ms	12,099ms
SQL Sever 2014	can unnest	cannot unnest

We cannot publish absolute runtime for SQL Server 2014, but you can guess from the asymptotics.

Unnesting is essential for good performance

- improves the asymptotics
- can lead to arbitrary gains

We present a generic approach for unnesting

- works on the algebra level, not on the SQL
- exploit set semantics, push down until no longer dependent
- can handle arbitrary queries
- virtually always beneficial, worst case memory overhead factor 2
- could often completely eliminate overhead, but that is a trade off