Updated 3/12/2022

Topic 3: Efficient query evaluation Unit 1: Acyclic query evaluation Lecture 16

Wolfgang Gatterbauer

CS7240 Principles of scalable data management (sp22)

https://northeastern-datalab.github.io/cs7240/sp22/ 3/11/2022

Islands of Tractability of CQ Evaluation

- Major Research Program: Identify <u>tractable cases</u> of the combined complexity of conjunctive query evaluation. $\partial = \frac{1}{2} + \frac{1}{2}$
- Over the years, this program has been pursued by two different research communities:
 - The Database Theory community
 - The Constraint Satisfaction community
- Explanation: Problems in those community are closely related:



Feder, Vardi: Monotone monadic SNP and constraint satisfaction, STOC 1993 <u>https://doi.org/10.1145/167088.167245</u> / Kolaitis, Vardi: Conjunctive-Query Containment and Constraint Satisfaction, JCSS 2000 <u>https://doi.org/10.1006/jcss.2000.1713</u> / Chandra, Merlin. "Optimal implementation of conjunctive queries in relational data bases", STOC 1977. <u>https://doi.org/10.1145/800105.803397</u> Based on Phokion Kolaitis' "Logic and Databases" series at Simons Institute, 2016. <u>https://simons.berkeley.edu/talks/logic-and-databases</u> Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u>

Acyclic queries

- Like many areas in computer science, cycles complicate things
 - Same with conjunctive query (CQ) evaluation
- Acyclic CQs are a large and useful tractable case CQs
- A query is acyclic if its relations can be placed in a tree (join tree) s.t.
 the set of nodes that contain any variable form a connected set
- Yannakakis' algorithm [81]: any acyclic query can be computed in time: O(|Input| + |Output|)

Yannakakis, "Algorithms for acyclic database schemes", VLDB 1981. <u>https://dl.acm.org/doi/10.5555/1286831.1286840</u> Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u>

Outline: T3-1: Acyclic conjunctive queries

- T3-1: Acyclic conjunctive queries
 - The semijoin operator
 - Join trees & Yannakakis algorithm
 - Query hypergraphs & GYO reduction
 - A detailed Yannakakis example
 - Full semijoin reductions
- T3-2: Cyclic conjunctive queries

Semijoin (⋉): derived RA operator



• Find tuples in **R** for which there is a matching tuple in **S** that is equal on their common attribute names.

$$\mathbf{R} \ltimes \mathbf{S} = \pi_{\mathbf{A}_1, \dots, \mathbf{A}_n} (\mathbf{R} \Join \mathbf{S})$$

where $A_1, ..., A_n$ are the attributes in R

Intuition: remove "dangling tuples" in R

• Example:

"Find actors who play some role."

RA:

Semijoin (⋉): derived RA operator



• Find tuples in **R** for which there is a matching tuple in **S** that is equal on their common attribute names.

SQL:

$$\mathbf{R} \ltimes \mathbf{S} = \pi_{\mathbf{A}_1, \dots, \mathbf{A}_n} (\mathbf{R} \Join \mathbf{S})$$

where $A_1, ..., A_n$ are the attributes in R

Intuition: remove "dangling tuples" in R

• Example:

"Find actors who play some role."

RA: Actor \ltimes Casts

Semijoin (⋉): derived RA operator



Actor (<u>aid</u>, name, gender) Play (<u>aid, mid, role</u>) Movie(<u>mid</u>, name, year)

• Find tuples in **R** for which there is a matching tuple in **S** that is equal on their common attribute names.



Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u>

Semijoins are often used to compute equijoins in distributed databases
 R \vee S
 Goal: send less data to reduce network bandwidth!



$$R \bowtie S = (R \ltimes S) \bowtie S$$
 law of semijoins

Task: compute with minimum data transfer:

Employee $\bowtie_{SSN=EmpSSN} (\sigma_{age>71} (Dependent)))$



Assumptions:

- 1. Very few employees have dependents.
- 2. Very few dependents have age > 71.
- 3. "Photo" is big.

Dependent			
EmpSSN	DepName	Age	Photo

Based upon an example by Dan Suciu Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u>

Task: compute with minimum data transfer:

Employee $\bowtie_{SSN=EmpSSN} (\sigma_{age>71} (Dependent)))$



Assumptions:

- 1. Very few employees have dependents.
- 2. Very few dependents have age > 71.
- 3. "Photo" is big.



Task: compute with minimum data transfer:

Employee $\bowtie_{SSN=EmpSSN} (\sigma_{age>71} (Dependent)))$



Assumptions:

- 1. Very few employees have dependents.
- 2. Very few dependents have age > 71.
- 3. "Photo" is big.



Based upon an example by Dan Suciu

Wolfgang Gatterbauer. Principles of scalable data management: https://northeastern-datalab.github.io/cs7240/

Task: compute with minimum data transfer:

Employee $\bowtie_{SSN=EmpSSN} (\sigma_{age>71} (Dependent)))$



Assumptions:

- 1. Very few employees have dependents.
- 2. Very few dependents have age > 71.
- 3. "Photo" is big.



Based upon an example by Dan Suciu

Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u>

Summary: Law of Semijoins

• Definition: the semi-join operation is:

$$\mathbf{R} \ltimes \mathbf{S} = \pi_{\mathbf{A}_1, \dots, \mathbf{A}_n} (\mathbf{R} \Join \mathbf{S})$$

- Formally, R ⋉ S means: retain from R only those tuples that have some matching tuple in S (in bag semantics: duplicates in R are preserved / Duplicates in S don't matter)
- Data complexity: O(|R| + |S|) ignoring log-factors
- Input: R(A₁,...,A_n), S(B₁,...,B_m), Output: T(A₁,...,A_n)
- The law of semijoins is:

 $R \bowtie S = (R \bowtie S) \bowtie S$

Thus, removing "dangling tuples" from a table does not change the query result



 $Q = (R \bowtie_B S \bowtie_C T)$

what is the problem?









 $Q = R \bowtie_B S \bowtie_C T$

what is the problem?

- Query output cardinality is 1
- But quadratic intermediate result sizes

Thus the query takes $O(n^2)$ despite constant output \circledast

SQL example available at: <u>https://github.com/northeastern-datalab/cs3200-activities/tree/master/sql</u> Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u>





$$Q = \mathbf{R} \Join_B \mathbf{S} \bowtie_C \mathbf{T}$$

what is the problem?

- Query output cardinality is 1
- But quadratic intermediate result sizes

Thus the query takes $O(n^2)$ despite constant output \otimes

what is a typical query plan?

 $Q = ((R \bowtie_B S) \bowtie_C T)$



What are the cardinalities at each stage?





$$Q = \mathbf{R} \Join_B \mathbf{S} \bowtie_C \mathbf{T}$$

what is the problem?

- Query output cardinality is 1
- But quadratic intermediate result sizes

Thus the query takes $O(n^2)$ despite constant output $\ensuremath{\mathfrak{O}}$

what is a typical query plan?



what are the cardinalities at each stage?





-- Query 1

select *

into record1

from R natural join S natural join T;

n=1000:

t_{Q1}=1451 msec

n=2000:

+q1=6104 msec

SQL example available at: <u>https://github.com/northeastern-datalab/cs3200-activities/tree/master/sql</u> Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u> ?

Can you think of a faster evaluation plan?

 t_{Q2} =5 msec t_{Q2} =8 msec





WITH clauses (also CTE = Common Table Expression) act like temporary views

-- Query 1

select *
into record1
from R natural join S natural join T;

n=1000:

t_{Q1}=1451 msec

n=2000:

+_{Q1}=6104 msec

SQL example available at: <u>https://github.com/northeastern-datalab/cs3200-activities/tree/master/sql</u> Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u>

-- Query 2 With S2 as (SELECT * FROM S WHERE S.B in (SELECT R.B FROM R)), S3 as (SELECT * FROM S2 WHERE S2.C in (SELECT T.C FROM T)) select a, b, c, d into record2 from R natural join S3 natural join T; $t_{Q2}=5$ msec $t_{Q2}=8$ msec

Diversion into CTE's (Common Table Expressions)





WITH clauses (also CTE = Common Table Expression) act like temporary views

- Allow to deconstruct more complex queries into simple blocks to be used and reused if necessary.
- Can increase readability by emphasizing a more <u>procedural</u> <u>interpretation</u> of a query in a workflow
- Especially useful if you need to reference a derived table multiple times in a single query or you perform <u>the same calculation</u> multiple times across multiple query components (= memoization)



The more general idea: "Sideways information passing"

Sideways information passing:

- "sending information from one subexpression not simply to its parent expression, but also to some other correlated portion of the query computation, in order to prune irrelevant results" [Ives, Taylor 08]
- includes techniques like two-way semijoins [Bernstein, Goodman 81] and magic sets [Beeri, Ramakrishnan 91]

$$\mathbf{Q} = \left((\mathbf{R} \bowtie_{\mathbf{B}} \mathbf{S}) \bowtie_{\mathbf{C}} \mathbf{T} \right)$$



[Bernstein, Goodman 81]. "Using Semi-Joins to Solve Relational Queries", JACM 1981. <u>https://doi.org/10.1145/322234.322238</u> [Beeri, Ramakrishnan 91]: "On the power of magic", Journal of Logic Programming, 1991. <u>https://doi.org/10.1016/0743-1066(91)90038-Q</u> Definition from: [Ives, Taylor 08]. "Sideways Information Passing for Push-Style Query Processing", ICDE 2008. <u>https://doi.org/10.1109/ICDE.2008.4497486</u> Wolfgang Gatterbauer. Principles of scalable data management: <u>https://northeastern-datalab.github.io/cs7240/</u>