# Query Optimization I 

CPS 216
Advanced Database Systems

## A query's trip through the DBMS



## Logical plan

* A tree whose nodes are logical operators
- Often a tree of relational algebra operators
- DB2 uses QGM (Query Graph Model)
* There are many equivalent logical plans



## Announcements (April 8)

Reading assignments for next week

- Selinger et al. "Access Path Selection in a Relational Database Management System." SIGMOD 1979
- Ioannidis and Kang. "Randomized Algorithms for Optimizing Large Join Queries." SIGMOD 1990
* Homework \#4 (short) will be assigned next

Tuesday and due the following Tuesday

* Final exam in 18 days (Monday, April 26)


## Parsing \& validation

* Parser: SQL $\rightarrow$ parse tree
- Good old lex \& yacc
- Detect and reject syntax errors
* Validator: parse tree $\rightarrow$ logical plan
- Detect and reject semantic errors
- Nonexistent tables/views/columns
- Type mismatches (e.g., AVG (name), name + GPA, Student UNION Enroll)
- Wildcard (SELECT *) and view expansion
- Use information stored in system catalog tables (contains all metadata/schema information)


## Query optimization

## * Conceptually

- Consider a space of possible plans (next)
- Estimate costs of plans in the search space (next Tuesday)
- Search through the space for the "best" plan (next Thursday)
* Often the goal is not picking the absolute optimum, but instead avoiding the horrible ones



## More relational algebra equivalences

* Convert $\sigma_{p}-\times$ to/from $\bowtie_{p}: \sigma_{p}(R \times S)=R \bowtie_{p} S$
* Merge/split $\sigma^{\prime}$ s: $\sigma_{p 1}\left(\sigma_{p 2} R\right)=\sigma_{p 1 \wedge p 2} R$
* Merge/split $\pi^{\prime}$ s: $\pi_{L 1}\left(\pi_{L 2} R\right)=\pi_{L 1} R$, where $L 1 \subseteq L 2$
* Push down/pull up $\sigma$ :
$\sigma_{p \wedge p r \wedge p s}(R \times S)=\left(\sigma_{p r} R\right) \bowtie_{p}\left(\sigma_{p s} S\right)$, where
- $p r$ is a predicate involving only $R$ columns
- $p s$ is a predicate involving only $S$ columns
- $p$ is a predicate involving both $R$ and $S$ columns
* Push down $\pi: \pi_{L}\left(\sigma_{p} R\right)=\pi_{L}\left(\sigma_{p}\left(\pi_{L L^{\prime}} R\right)\right)$, where
- $L^{\prime}$ is the set of columns referenced by $p$ that are not in $L$
* Many more (seemingly trivial) equivalences...
- Can be systematically used to transform a plan to new ones


## Heuristics-based query optimization

* Start with a logical plan
* Push selections/projections down as much as possible
- Why? Reduce the size of intermediate results
- Why not? May be expensive; maybe joins filter better
* Join smaller relations first, and avoid cross product
- Why? Reduce the size of intermediate results
- Why not? Size depends on join selectivity too
* Convert the transformed logical plan to a physical plan (by choosing appropriate physical operators)


## Plan enumeration in relational algebra

* Apply relational algebra equivalences
$\infty$ Join reordering: $\times$ and $\bowtie$ are associative and commutative (except when column ordering is considered, but that is unimportant)



## Relational query rewrite example



## SQL query rewrite

- More complicated—subqueries and views divide a query into nested "blocks"
- Processing each block separately forces particular join methods and join order
- Even if the plan is optimal for each block, it may not be optimal for the entire query
* Unnest query: convert subqueries/views to joins
$\sigma$ Then we just deal with select-project-join queries
- Where the clean rules of relational algebra apply


## DB2's QGM

Leung et al. "Query Rewrite Optimization Rules in IBM DB2 Universal Database."

* Query Graph Model: DB2's logical plan language
- More high-level than relational algebra
* A graph of boxes
- Leaf boxes are tables
- The standard box is the SELECT box (actually a select-project-join query block with optional duplicate elimination)
- Other types include GROUPBY (aggregation), UNION, INTERSECT, EXCEPT
- Can always add new types (e.g., OUTERJOIN)



## E to F conversion

* SELECT DISTINCT name FROM Student WHERE SID = ANY (SELECT SID FROM Enroll);
* SELECT DISTINCT name

FROM Student, (SELECT SID FROM Enroll) t WHERE Student.SID = t.SID;
(EtoF rule)

* SELECT DISTINCT name

FROM Student, Enroll
WHERE Student.SID = Enroll.SID;
(SELMERGE rule)

## More on QGM boxes

* Head: declarative description of the output
- Schema: list of output columns
- Property: Are output tuples DISTINCT?
* Body: how to compute the output
- Quantifiers: tuple variables that range over other boxes
- F: regular tuple variable, e.g., FROM $R$ AS $r$
- E: existential quantifier, e.g., IN (subquery), or = ANY (subquery)
- A: universal quantifier, e.g., > ALL (subquery)
- S: scalar subquery, e.g., $=($ subquery $)$
- Quantifiers are connected a hypergraph
- Hyperedges are predicates
- Enforce DISTINCT, preserve duplicates, or permit duplicates?
- For the output of this box, and for each quantifier


## Query rewrite in DB2

* Goal: make the logical plan as general as possible, i.e., merge boxes
* Rule-based transformations on QGM
- Merge subqueries in FROM
- Convert E to F (e.g., IN/ANY subqueries to joins)
- Convert intersection to join
- Convert S to F (i.e., scalar subqueries to joins)
- Convert outerjoin to join
- Magic (i.e., correlated subqueries to joins)


## Problem with duplicates

Same query, without DISTINCT

* SELECT name

FROM Student
WHERE SID = ANY (SELECT SID FROM Enroll);

* SELECT name

FROM Student, Enroll WHERE Student.SID = Enroll.SID;

* Suppose some student takes multiple classes
- The first query returns name once; the second multiple times
* Adding DISTINCT to the second query does not help
- Suppose two students have the same name


## A way of preserving duplicates

```
* SELECT name
    FROM Student
    WHERE SID = ANY (SELECT SID FROM Enroll);
```

Suppose that SID is a key of Student

* SELECT DISTINCT Student.SID, name

FROM Student, Enroll
WHERE Student.SID = Enroll.SID;
(ADDKEYS rule)

* Then simply project out Student. SID


## Another E to F trick

* Sometimes an ANY subquery can be turned into an aggregate subquery without ANY, to improve performance further
* SELECT * FROM Student s1 WHERE GPA > ANY
(SELECT GPA FROM Student s2 WHERE s2.name = 'Bart');
* SELECT * FROM Student s 1 WHERE GPA > (SELECT MIN(GPA) FROM Student s2 WHERE s2.name = 'Bart');


## Does the same trick apply to ALL?

* SELECT * FROM Student s1 WHERE GPA > ALL
(SELECT GPA FROM Student s2
WHERE s2.name = 'Bart');
* SELECT * FROM Student s1

WHERE GPA >
(SELECT MAX (GPA) FROM Student s2
WHERE s2.name = 'Bart');

* Suppose there is no student named Bart
- The first query returns all students; the second returns none


## COUNT bug

* SELECT CID FROM Course

WHERE title LIKE 'CPS\%'
AND min_enroll > (SELECT COUNT (*) FROM Enroll WHERE Enroll.CID = Course.CID);

* SELECT CID First compute the enrollment for all(?) courses FROM Course, (SELECT CID, COUNT (*) AS cnt

FROM Enroll GROUP BY CID) $t$
WHERE t.CID = Course.CID AND min_enroll > t.cnt AND title LIKE 'CPS\%';

* Suppose a CPS class is empty
- The first query returns this course; the second does not


## Magic decorrelation example

* SELECT CID FROM Course

WHERE title LIKE 'CPS\%'
AND min_enroll > (SELECT COUNT (*) FROM Enroll WHERE Enroll.CID = Course.CID);

* CREATE VIEW Supp_Course AS Process the outer query SELECT * FROM Coūrse WHERE title LIKE 'CPS\%'; without the subquery
CREATE VIEW Magic AS
SELECT DISTINCT CID FROM Supp_Course;
Collect bindings
CREATE VIEW DS AS
Evaluate the subquery
(SELECT Enrol1.CID, COUNT (*) AS cnt with bindings FROM Magic, Enroll WHERE Magic.CID = Enroll.CID GROUP BY Enroll.CID) UNION
(SELECT Magic.CID, 0 AS cnt FROM Magic
WHERE Magic.CID NOT IN (SELECT CID FROM Enroll);
SELECT Supp_Course.CID FROM Supp_Course, DS Finally, refine WHERE Supp_Course.CID = DS.CID $\quad$ the outer query AND min_enroll > DS.cnt;


## Summary of query rewrite

* Break the artificial boundary between queries and subqueries
* Combine as many query blocks as possible in a select-project-join block, where the clean rules of relational algebra apply
* Handle with care-extremely tricky with duplicates, NULL's, empty tables, and correlation

