Evaluation of Relational Operations: Other Techniques

Chapter 14

Sayyed Nezhadi

Schema for Examples

Sailors (*sid*: integer, *sname*: string, *rating*: integer, *age*: real)
Reserves (*sid*: integer, *bid*: integer, *day*: dates, *rname*: string)

- Reserves:
  - Each tuple is 40 bytes long, 100 tuples per page, 1000 pages.

- Sailors:
  - Each tuple is 50 bytes long, 80 tuples per page, 500 pages.

Last Week

- Index nested loop

- Sort-merge join

Simple Nested Loops Join

foreach tuple r in R do
    foreach tuple s in S do
        if r_i == s_j then add <r, s> to result

- For each tuple in the *outer* relation R, we scan the entire *inner* relation S.
  - Cost: $M + p_R * M * N$ = 1000 + 100*1000*500 I/Os.
- Page-oriented Nested Loops join: For each *page* of R, get each *page* of S, and write out matching pairs of tuples   <r, s>, where r is in R-page and S is in S-page.
  - Cost: $M + M*N$ = 1000 + 1000*500
  - If smaller relation (S) is outer, cost = 500 + 500*1000

Block Nested Loops Join

- Use one page as an input buffer for scanning the inner S, one page as the output buffer, and use all remaining pages to hold ``block'' of outer R.
  - For each matching tuple r in R-block, s in S-page, add <r, s> to result. Then read next R-block, scan S, etc.

R & S

Hash table for block of R
(k < B-1 pages)

Join Result
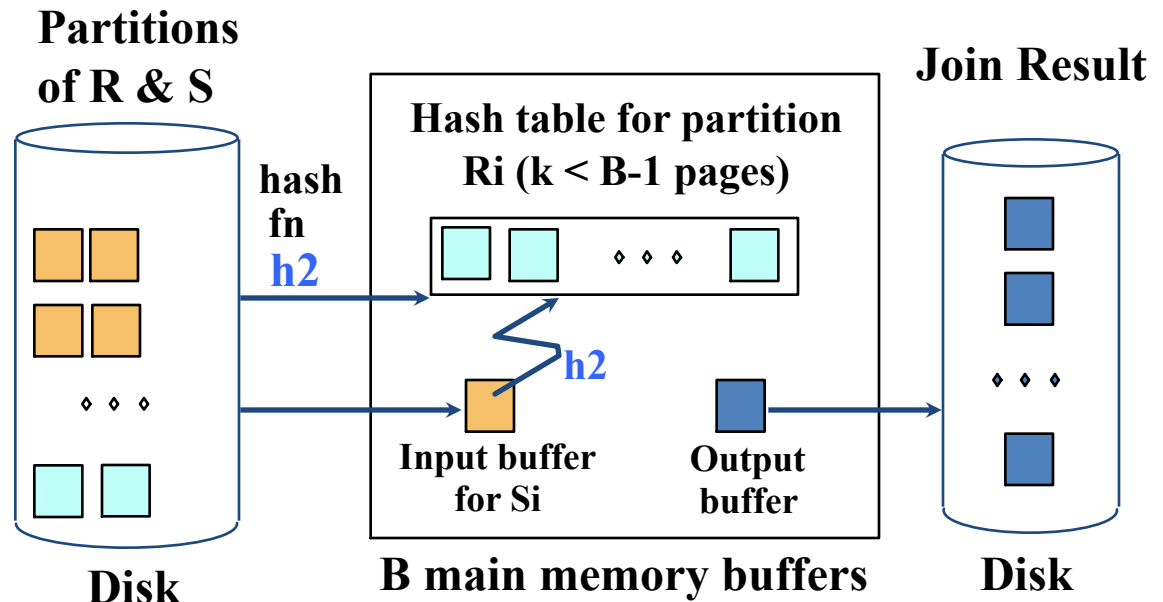
Input buffer for S   Output buffer
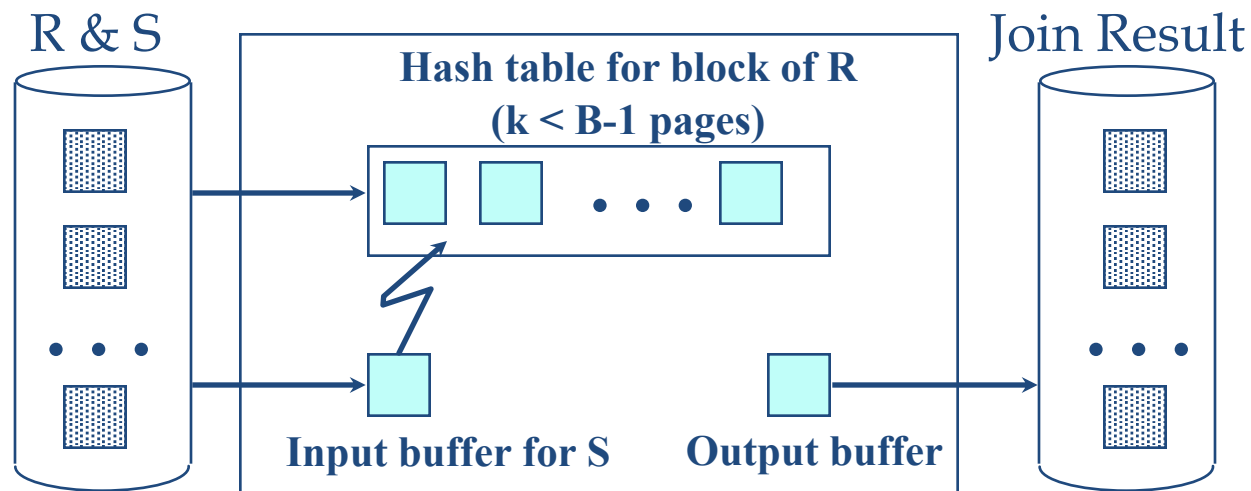
Examples of Block Nested Loops

- Cost: Scan of outer + #outer blocks * scan of inner
  - #outer blocks = $\lceil \# \text{ of pages of outer } / \text{ blocksize} \rceil$
- With Reserves (R) as outer, and 100 pages of R:
  - Cost of scanning R is 1000 I/Os; a total of 10 *blocks*.
  - Per block of R, we scan Sailors (S); 10*500 I/Os.
  - If space for just 90 pages of R, we would scan S 12 times.
- With 100-page block of Sailors as outer:
  - Cost of scanning S is 500 I/Os; a total of 5 blocks.
  - Per block of S, we scan Reserves;   5*1000 I/Os.
- With *sequential reads* considered, analysis changes: may be best to divide buffers evenly between R and S.

Hash-Join

- Partition both relations using hash fn **h**: R tuples in partition i will only match S tuples in partition i.

- ❖ Read in a partition of R, hash it using **h2 (<> h!)**. Scan matching partition of S, search for matches.

Original Relation

OUTPUT

Partitions

INPUT

hash function
h

1

2

B-1

1

2

B-1

Disk

B main memory buffers

Disk

Partitions of R & S

Join Result

hash fn
h2

Hash table for partition Ri (k < B-1 pages)

h2

Input buffer for Si

Output buffer

Disk

B main memory buffers

Disk

Observations on Hash-Join

- #partitions k <= B-1 (why?), and B-2 > size of largest partition to be held in memory. Assuming uniformly sized partitions, and maximizing k, we get:
  - k= B-1, and M/(B-1) < B-2, i.e., B must be > $\sqrt{M}$
- If we build an in-memory hash table to speed up the matching of tuples, a little more memory is needed.
- If the hash function does not partition uniformly, one or more R partitions may not fit in memory. Can apply hash-join technique recursively to do the join of this R-partition with corresponding S-partition.

Cost of Hash-Join

- In partitioning phase, read+write both relns; 2(M+N). In matching phase, read both relns; M+N I/Os.
- In our running example, this is a total of 4500 I/Os.
- Sort-Merge Join vs. Hash Join:
    - Given a minimum amount of memory (*what is this, for each?*) both have a cost of 3(M+N) I/Os. Hash Join superior on this count if relation sizes differ greatly. Also, Hash Join shown to be highly parallelizable.
    - Sort-Merge less sensitive to data skew; result is sorted.

General Join Conditions

- Equalities over several attributes (e.g., *R.sid=S.sid* AND *R.rname=S.sname*):
    - For Index NL, build index on *<sid, sname>* (if S is inner); or use existing indexes on *sid* or *sname.*
    - For Sort-Merge and Hash Join, sort/partition on combination of the two join columns.

- Inequality conditions (e.g., *R.rname < S.sname*):
    - For Index NL, need (clustered!) B+ tree index.
        - Range probes on inner; # matches likely to be much higher than for equality joins.
    - Hash Join, Sort Merge Join not applicable.
    - Block NL quite likely to be the best join method here.

Using an Index for Selections

- Cost depends on #qualifying tuples, and clustering.
  - Cost of finding qualifying data entries (typically small) plus cost of retrieving records (could be large w/o clustering).
  - In example, assuming uniform distribution of names, about 10% of tuples qualify (100 pages, 10000 tuples). With a clustered index, cost is little more than 100 I/Os; if unclustered, upto 10000 I/Os!
- *Important refinement for unclustered indexes*:
    1. Find qualifying data entries.
    2. Sort the rid's of the data records to be retrieved.
    3. Fetch rids in order. This ensures that each data page is looked at just once (though # of such pages likely to be higher than with clustering).

Two Approaches to General Selections

- First approach: Find the *most selective access path*, retrieve tuples using it, and apply any remaining terms that don't match the index:

  - *Most selective access path:* An index or file scan that we estimate will require the fewest page I/Os.

  - Terms that match this index reduce the number of tuples *retrieved*; other terms are used to discard some retrieved tuples, but do not affect number of tuples/pages fetched.

  - Consider *day<8/9/94 AND bid=5 AND sid=3.* A B+ tree index on *day* can be used; then, *bid=5* and *sid=3* must be checked for each retrieved tuple. Similarly, a hash index on *<bid, sid>* could be used; *day<8/9/94* must then be checked.

Intersection of Rids

- <u>Second approach</u> (if we have 2 or more matching indexes that use Alternatives (2) or (3) for data entries):
  - Get sets of rids of data records using each matching index.
  - Then *intersect* these sets of rids (we'll discuss intersection soon!)
  - Retrieve the records and apply any remaining terms.
  - Consider *day<8/9/94 AND bid=5 AND sid=3.* If we have a B+ tree index on *day* and an index on *sid*, both using Alternative (2), we can retrieve rids of records satisfying *day<8/9/94* using the first, rids of recs satisfying *sid=3* using the second, intersect, retrieve records and check *bid=5.*

The Projection Operation

- An approach based on sorting:
  - Modify Pass 0 of external sort to eliminate unwanted fields. Thus, runs of about 2B pages are produced, but tuples in runs are smaller than input tuples. (Size ratio depends on # and size of fields that are dropped.)
  - Modify Pass 0 & merging passes to eliminate duplicates. Thus, number of result tuples smaller than input. (Difference depends on # of duplicates.)
  - Cost: In Pass 0, read original relation (size M), write out same number of smaller (distinct) tuples. In merging passes, fewer tuples written out in each pass. Reserves: 1000 input pages reduced to 250 in Pass 0 if size ratio is 0.25

Projection Based on Hashing

- *Partitioning phase*: Read R using one input buffer. For each tuple, discard unwanted fields, apply hash function *h1* to choose one of B-1 output buffers.
  - Result is B-1 partitions (of tuples with no unwanted fields). 2 tuples from different partitions guaranteed to be distinct.
- *Duplicate elimination phase*: For each partition, read it and build an in-memory hash table, using hash fn *h2* (<> *h1*) on all fields, while discarding duplicates.
  - If partition does not fit in memory, can apply hash-based projection algorithm recursively to this partition.
- Cost: For partitioning, read R, write out each tuple, but with fewer fields. This is read in next phase.

Discussion of Projection

- Sort-based approach is the standard; better handling of skew and result is sorted.
- If an index on the relation contains all wanted attributes in its search key, can do *index-only* scan.
  - Apply projection techniques to data entries (much smaller!)
- If an ordered (i.e., tree) index contains all wanted attributes as *prefix* of search key, can do even better:
  - Retrieve data entries in order (index-only scan), discard unwanted fields, compare adjacent tuples to check for duplicates.

Set Operations

- Intersection and cross-product special cases of join.
- Union (Distinct) and Except similar; we'll do union.
- Sorting based approach to union:
  - Sort both relations (on combination of all attributes).
    - Remove duplicates? (if not base relations)
  - Scan sorted relations and merge them.
  - *Alternative*: Merge runs from Pass 0 for *both* relations.
- Hash based approach to union:
  - Partition R and S using hash function h.
  - For each S-partition, build in-memory hash table (using $h2$), scan corresponding R-partition and add tuples to table while discarding duplicates.

Aggregate Operations (AVG, MIN, etc.)

- Without grouping:
  - In general, requires scanning the relation.
  - Given index whose search key includes all attributes in the SELECT (if there is no WHERE), can do index-only scan.
- With grouping:
  - Sort on group-by attributes, then scan relation and compute aggregate for each group. (Can improve upon this by combining sorting and aggregate computation.)
  - Similar approach based on hashing on group-by attributes.
  - Given tree index whose search key includes all attributes in SELECT, WHERE and GROUP BY clauses, can do index-only scan; if group-by attributes form prefix of search key, can retrieve data entries/tuples in group-by order.

Impact of Buffering

- If several operations are executing concurrently, estimating the number of available buffer pages is guesswork.

- Repeated access patterns interact with buffer replacement policy.
  - Nested Join
  - e.g., Inner relation is scanned repeatedly in Simple Nested Loop Join. With enough buffer pages to hold inner, replacement policy does not matter. Otherwise, MRU is best, LRU is worst (*sequential flooding*).
  - Does replacement policy matter for Block Nested Loops?
  - What about Index Nested Loops? Sort-Merge Join?

Highlights of System R Optimizer

- Impact:
  - Most widely used currently; works well for < 10 joins.
- Cost estimation: Approximate art at best.
  - Statistics, maintained in system catalogs, used to estimate cost of operations and result sizes.
  - Considers combination of CPU and I/O costs.
- Plan Space: Too large, must be pruned.
  - Only the space of *left-deep plans* is considered.
    - Left-deep plans allow output of each operator to be *pipelined* into the next operator without storing it in a temporary relation.
  - Cartesian products avoided.

Cost Estimation

- For each plan considered, must estimate cost:
  - Must estimate *cost* of each operation in plan tree.
    - Depends on input cardinalities.
    - We've already discussed how to estimate the cost of operations (sequential scan, index scan, joins, etc.)
  - Must also estimate *size of result* for each operation in tree!
    - Use information about the input relations.
    - For selections and joins, assume independence of predicates.
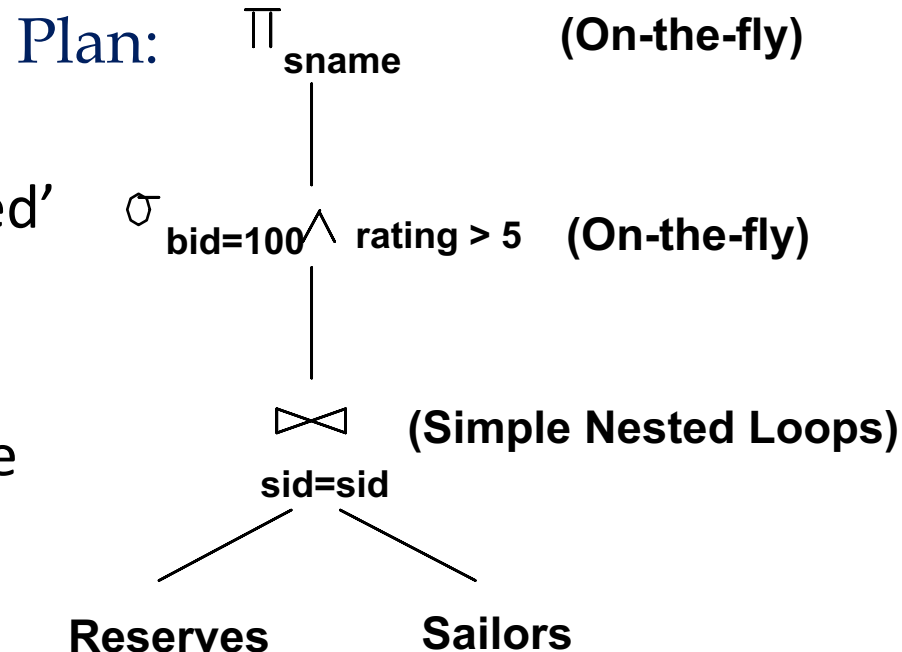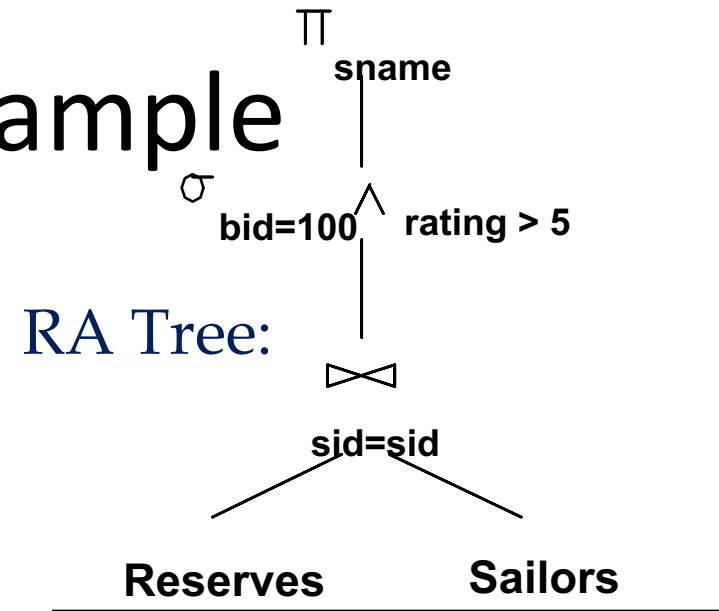
Size Estimation and Reduction Factors

- Consider a query block:

> SELECT attribute list
> FROM relation list
> WHERE term1 AND … AND termk

- Maximum # tuples in result is the product of the cardinalities of relations in the FROM clause.

- *Reduction factor (RF)* associated with each *term* reflects the impact of the *term* in reducing result size. *Result cardinality* = Max # tuples * product of all RF's.

  – Implicit assumption that *terms* are independent!

  – Term *col=value* has RF *1/NKeys(I),* given index I on *col*

  – Term *col1=col2* has RF *1/MAX(NKeys(I1), NKeys(I2))*

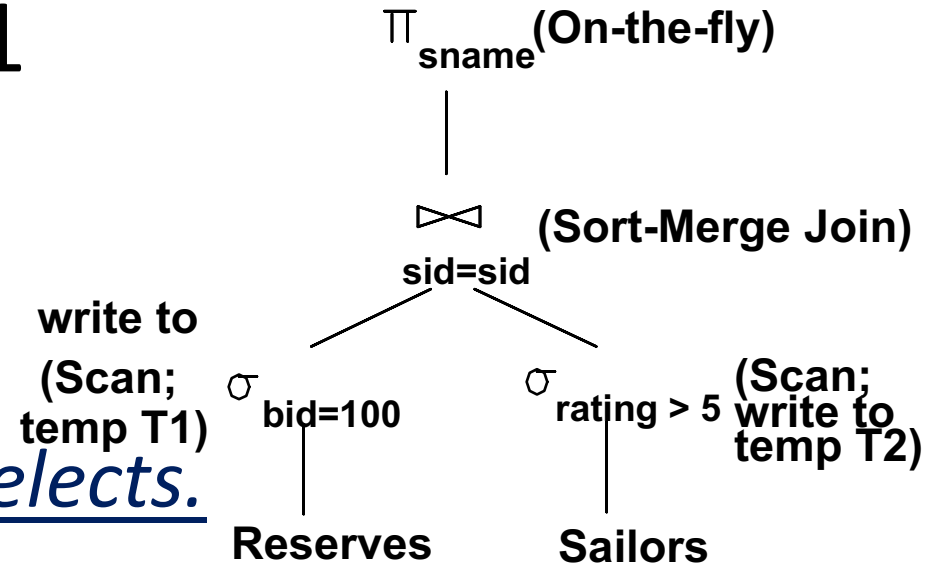 – Term *col>value* has RF *(High(I)-value)/(High(I)-Low(I))*

Motivating Example

Π_{sname}

SELECT S.sname
FROM Reserves R, Sailors S
WHERE R.sid=S.sid AND
    R.bid=100 AND S.rating>5

RA Tree:

σ bid=100 $\wedge$ rating > 5

$\bowtie$ sid=sid

Reserves   Sailors

- Cost: 500+500*1000 I/Os

- By no means the worst plan!

- Misses several opportunities: selections could have been `pushed' earlier, no use is made of any available indexes, etc.

- *Goal of optimization:* To find more efficient plans that compute the same answer.

Plan: Π_{sname} (On-the-fly)

σ bid=100 $\wedge$ rating > 5   (On-the-fly)

$\bowtie$ (Simple Nested Loops)
sid=sid

Reserves   Sailors

Alternative Plans 1 (No Indexes)

π_{sname} **(On-the-fly)**

$\bowtie$ **(Sort-Merge Join)**
sid=sid

write to
(Scan; $\sigma_{bid=100}$ $\sigma_{rating > 5}$ **(Scan;**
temp T1)   **write to**
                                      temp T2)

Reserves   **Sailors**

- ***Main difference: <u>push selects.</u>***

- With 5 buffers, cost of plan:
  - Scan Reserves (1000) + write temp T1 (10 pages, if we have 100 boats, uniform distribution).
  - Scan Sailors (500) + write temp T2 (250 pages, if we have 10 ratings).
  - Sort T1 (2*2*10), sort T2 (2*4*250), merge (10+250)
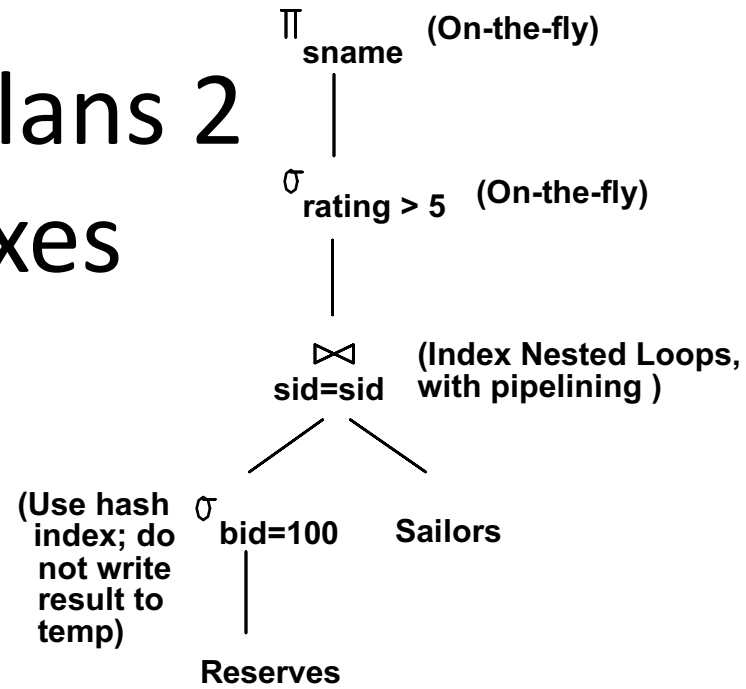 - Total: 4060 page I/Os.

Buffered Nested Loop(BNL)

- Buffered Nested Loop
  - Bring pages to memory in Groups
  - Hash the pages
  - Scan the second table to find matches
  - For each 3 page block from T1, scan the entire T2
- BNL, join cost = 10+4*250, total cost = 2770, why?
- If we `push' projections, T1 has only *sid*, T2 only *sid* and *sname*:
  - T1 fits in 3 pages, cost of BNL drops to under 250 pages, total < 2000.

Alternative Plans 2 With Indexes

Π_{sname} **(On-the-fly)**

$\sigma_{\text{rating > 5}}$ **(On-the-fly)**

$\bowtie_{\text{sid=sid}}$ **(Index Nested Loops, with pipelining)**

(Use hash index; do not write result to temp) $\sigma_{\text{bid=100}}$ **Sailors**

Reserves

- With clustered index on *bid* of Reserves, we get 100,000/100 = 1000 tuples on 1000/100 = 10 pages.

- INL with ***pipelining*** (outer is not materialized).

    – Projecting out unnecessary fields from outer doesn't help.

❖ Join column *sid* is a key for Sailors.

    – At most one matching tuple, unclustered index on *sid* OK.

❖ Decision not to push *rating>5* before the join is based on availability of *sid* index on Sailors.

❖ Cost: Selection of Reserves tuples (10 I/Os); for each, must get matching Sailors tuple (1000*1.2); total 1210 I/Os.

Summary

- There are several alternative evaluation algorithms for each relational operator.
- A query is evaluated by converting it to a tree of operators and evaluating the operators in the tree.
- Must understand query optimization in order to fully understand the performance impact of a given database design (relations, indexes) on a workload (set of queries).
- Two parts to optimizing a query:
  – Consider a set of alternative plans.
    - Must prune search space; typically, left-deep plans only.
  – Must estimate cost of each plan that is considered.
    - Must estimate size of result and cost for each plan node.
    - *Key issues*: Statistics, indexes, operator implementations.