Overview

Join algorithms

- (Block) Nested Loop Join - two-level for-loop
- Hash Join - compute a hash table of one input; probe the hash table with the other input
- Sort-Merge Join - sort both tables on one of the join conditions, then merge sorted lists

Indexes

- B-Tree index - supports point and range lookup.
- Hashtable - supports point lookup.

Clustered indexes are the way the data is stored in the “original” table; they are almost always B-trees. Unclustered indexes define secondary tables that reference the main table via pointers.

Cardinality estimation - the problem of estimating the number of tuples after an operation, such as a selection or join. Good estimates are critical to cost modeling; the larger the cardinality; the larger the cost. Here we use assumptions: that the values of a table are uniformly distributed among its distinct values, and that all joins are foreign key-primary key joins.

Cost Modeling

- B(R) - the number of blocks used to store the relation R on disk
- T(R) - the number of tuples in R (also known as R's cardinality)
- V(R, a) - the number of unique values of attribute a in relation R
- M - the number of pages that fit in memory

Cost-based Query Optimization compares plans by computing their estimated cost, then chooses the one with the cheapest estimated cost to execute.

Query execution - we learned about the iterator method (iterator interface) of executing the operators in a query plan. You may see it called the "pull-based model of query execution", because each operator "pulls" data from its child operators by calling next(). The three methods used are open(), next(), and close().

Problems

1. (Adapted from 414 SP 17 Final)

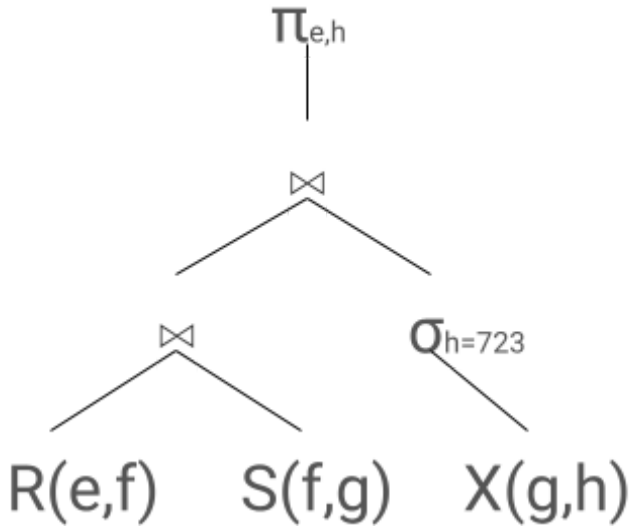

Consider the relations R(e, f), S(f, g), and X(g, h) in the query plan depicted above.

- Joins are natural joins.
- Every attribute is integer-valued.
- Assume that **every intermediate result is materialized** (i.e., written to disk).
- Assume that we are executing queries on a machine that has **11 memory pages** available.
- Assume **uniform distributions** on the attributes for the purpose of computing estimates.

Consider the following statistics:

Table	#tuples	#blocks
R	1,000	100
S	5,000	200
X	100,000	10,000

Attribute	# distinct values	Minimum value	Maximum value
R.f	100	1	1,000
S.f	1,000	1	2,000
S.g	5,000	1	2,000
X.g	1,000	1	10,000
X.h	1,000	1	500,000

A. Estimate the number of tuples and blocks in the selection $\sigma_{h=723}(X)$.
Select * from X where h = 723

B. Estimate the number of tuples in the join $R \bowtie S$.

2. Assume you have two relations, R(a,b) and S(b,c) with the following statistics:

B(R) = 4000 T(R) = 6×10^5 V(R,A) = 2×10^3 V(R,B) = 5×10^4	B(S) = 9×10^4 T(S) = 3×10^6 V(S,B) = 3×10^4 V(S,C) = 10^3

Additionally,

- Assume that indexes and intermediate results fit in memory, using pipelined execution.
- There are clustered indexes on R.a and S.b
- There are unclustered indexes on R.b and S.c
- Assume that 123 is a value of R.a and 456 is a value of S.c; that the values are evenly distributed

Two equivalent physical query plans are drawn below. Calculate the estimated IO cost of each operation lettered in the plans. Your answers should be integers, but show work for credit.

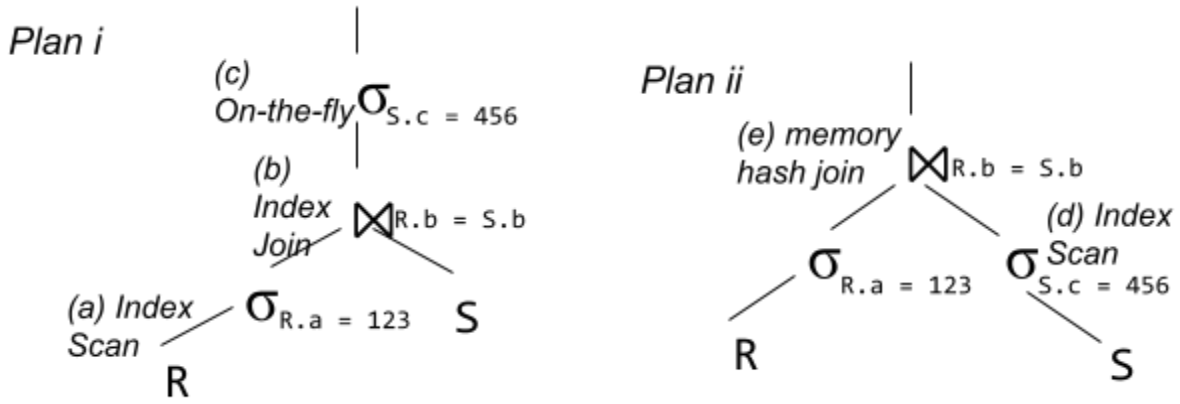

a) (2 points)_____

b) (3 points) _____

c) (2 points) _____

d) (2 points) _____

e) (3 points) _____

f) (3 points) What is the total IO cost of each plan? Which plan would you choose to minimize the total IO cost?