

1. **Join Cost Calculation:**
Given two relations $R(A, B)$ with 10,000 tuples and $S(B, C)$ with 50,000 tuples, and memory can hold 100 tuples at a time, calculate the cost of a nested loop join.
2. **Indexing Cost:**
A B+ tree index is built on a table with 1 million records, and the tree has 4 levels. How many disk accesses are required to retrieve a record?
3. **Query Execution Time:**
If a full table scan takes 500ms and an index scan takes 50ms, how much faster is the index scan?
4. **Cost of Sorting:**
If sorting a table of 1 million rows takes $O(n \log n)$ time, estimate the time required if $n = 10^6$.
5. **Selection Query Cost:**
A table has 5 million records, and 10% of the records satisfy a condition. If a sequential scan takes 0.1ms per record, how long does the query take?
6. **Partitioning Benefit Calculation:**
A table has 10 million rows, partitioned into 10 partitions. If a query searches for records in one partition, how much data scanning is reduced?
7. **Materialized View Storage Cost:**
A materialized view stores 1GB of precomputed query results, but recalculating the query from the base table takes 500MB of temporary storage. How much extra storage is used?
8. **I/O Cost for Index Lookup:**
If an indexed lookup takes 3 disk accesses and each disk access takes 5ms, what is the total retrieval time?
9. **Query Optimization Decision:** Given three query plans.
Plan A: 200ms, Plan B: 150ms, Plan C: 120ms. Which plan should be selected and why?
10. What are the advantages of using Indexes in Query Evaluation?
11. Explain the differences between Nested Loop Join, and Sort-Merge Join.
12. What is the role of Query Execution Plans in DBMS? How can we analyze them using EXPLAIN ANALYZE?