CARNEGIE MELLON UNIVERSITY
COMPUTER SCIENCE DEPARTMENT
15-445/645 – DATABASE SYSTEMS (FALL 2026)
PROF. JIGNESH PATEL

Homework #4 (by Drew)
Due: **Sunday, November 1, 2026 @ 11:59pm**

IMPORTANT:

- Enter all of your answers into **Gradescope by 11:59pm on Sunday, November 1, 2026**.
- **Plagiarism**: Homework may be discussed with other students, but all homework is to be completed **individually**.

For your information:

- Graded out of **100** points; **3** questions total
- Rough time estimate: $\approx$ 2 - 3 hours (0.5 - 1 hours for each question)

Revision : 2026/10/23 02:33

Question	Points	Score
Sorting Algorithms	32	
Join Algorithms	44	
Query Execution, Planning, and Optimization	24	
Total:	100	

Question 1: Sorting Algorithms . . . . . [32 points]

We have a database file with 12 million pages (N = 12,000,000 pages), and we want to sort it using external merge sort. Assume that the DBMS is not using double buffering or blocked I/O, and that it uses quicksort for in-memory sorting. Let B denote the number of buffers.

(a) **[4 points]** Assume that the DBMS has 120 buffers. How many total sorted runs are generated across all passes? Note that the final sorted file does not count towards the sorted run count.
☐ 34 ☐ 80,000 ☐ 81,667 ☐ 100,849 ☐ 100,850 ☐ 10,848

(b) **[4 points]** Again, assuming that the DBMS has 120 buffers. How many passes does the DBMS need to perform in order to sort the file?
☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6

(c) **[4 points]** Again, assuming that the DBMS has 120 buffers. How many pages does each sorted run have after the third pass (i.e. Note: this is Pass #2 if you start counting from Pass #0)?
☐ 119 ☐ 120 ☐ 640,399 ☐ 640,400 ☐ 640,401 ☐ 1,699,320
☐ 39,443,120

(d) **[4 points]** Again, assuming that the DBMS has 120 buffers. What is the total I/O cost to sort the file?
☐ 8,000,000 ☐ 32,000,000 ☐ 64,000,000 ☐ 96,000,000 ☐ 192,000,000

(e) **[4 points]** What is the smallest number of buffers B such that the DBMS can sort the target file using only three passes?
☐ 106 ☐ 230 ☐ 160 ☐ 240 ☐ 241

(f) **[4 points]** Suppose the DBMS has 67 buffers. What is the largest database file (expressed in terms of the number of pages) that can be sorted with external merge sort using three passes?
☐ 11,342 ☐ 120,050 ☐ 278,852 ☐ 291,852 ☐ 300,763

(g) i. **[4 points]** The DBMS receives a query that requires sorting. Assume that the sort order is a prefix of the index key. Under which scenario(s) will using an unclustered B+Tree index have comparable performance to a clustered B+Tree index:
☐ The sort order exactly matches the index key.
☐ All attributes accessed by the query are contained in the index.
☐ Query contains a LIMIT 1, and the first tuple answers the query.
☐ Using unclustered index will always perform worse than using clustered index.
☐ None of the above.

ii. **[2 points]** It is always less efficient for a DBMS to use a sort aggregate than a hash aggregate.
☐ True ☐ False

Question 1 continues. . .

iii. **[2 points]** Hash aggregation **must** be used for some of the aggregates discussed in lecture (i.e., "DISTINCT", "GROUP BY").
☐ True ☐ False

Question 2: Join Algorithms . . . . . [44 points]

Consider relations `X(a, b)`, `Y(a, c, e)`, and `Z(a, d, f)` to be joined on the common attribute `a`. Assume that there are no indexes available on the tables to speed up the join algorithms.

- There are B = 100 pages in the buffer
- Table `X` spans M = 1,000 pages with 100 tuples per page
- Table `Y` spans N = 500 pages with 500 tuples per page
- Table `Z` spans O = 3,000 pages with 250 tuples per page
- The join result of `Y` and `Z` spans P = 50 pages

For the following questions, assume a simple cost model where pages are read and written one at a time. Also assume that one buffer block is needed for the evolving output block and one input block is needed for the current input block of the inner relation. You may ignore the cost of the writing of the final results.

(a) **[2 points]** What is the I/O cost of a simple nested loop join with `Y` as the outer relation and `X` as the inner relation?

- ☐ 400,500
- ☐ 500,500
- ☐ 80,000,500
- ☐ 200,000,500
- ☐ 250,000,500
- ☐ 40,000,000,500

(b) **[2 points]** What is the I/O cost of a block nested loop join with `Y` as the outer relation and `Z` as the inner relation?

- ☐ 6,500
- ☐ 9,000
- ☐ 9,500
- ☐ 18,500
- ☐ 1,500,000
- ☐ 1,500,500

(c) **[2 points]** What is the I/O cost of a block nested loop join with `Z` as the outer relation and `Y` as the inner relation?

- ☐ 8,000
- ☐ 10,576
- ☐ 18,500
- ☐ 19,500
- ☐ 1,500,000
- ☐ 1,503,000

(d) For a sort-merge join with `Z` as the outer relation and `X` as the inner relation:

i. **[3 points]** What is the cost of sorting the tuples in `X` on attribute `a`?
- ☐ 4,000
- ☐ 8,000
- ☐ 9,200
- ☐ 12,000
- ☐ 16,000

ii. **[3 points]** What is the cost of sorting the tuples in `Z` on attribute `a`?
- ☐ 8,000
- ☐ 9,000
- ☐ 11,000
- ☐ 12,000
- ☐ 24,000

iii. **[3 points]** What is the cost of the merge phase in the worst-case scenario?
- ☐ 2,800,000
- ☐ 2,800,000
- ☐ 2,870,000
- ☐ 3,000,000
- ☐ 90,000,000,000

iv. **[3 points]** What is the cost of the merge phase assuming there are no duplicates in the join attribute?
- ☐ 340
- ☐ 1,640
- ☐ 2,000
- ☐ 3,000
- ☐ 4,000
- ☐ 28,000

v. **[3 points]** Now consider joining `Y`, `Z` and then joining the result with `X`. What is the cost of the final merge phase assuming there are no duplicates in the join attribute?
- ☐ 420
- ☐ 1,000
- ☐ 1,050
- ☐ 4,000
- ☐ 22,000

(e) **[2 points]** Consider a partitioned hash join with `Y` as the outer relation and `X` as the inner relation. You may ignore recursive partitioning and partially filled blocks. What is the cost of the combined probe and partition phases?
- ☐ 1,500
- ☐ 3,000

□ 4,225
□ 4,500
□ 9,000

(f) **[3 points]** Assume that the tables do not fit in main memory and that a large number of distinct values hash to the same bucket using hash function h_1. Which of the following approaches works the best?
□ Create hashtables for the inner and outer relation using h_1 and rehash into an embedded hash table using a second different hash function, h_2, for large buckets.
□ Build a bloom filter over the inner relation with h_1 to test for membership before checking the inner relation's hashtable.
□ Use linear probing for collisions and page in and out parts of the hashtable needed at a given time.
□ Create hashtables for the inner and outer relation using h_1 and rehash into an embedded hash table using h_1 for large buckets.

(g) For each of the following statements about joins, pick True or False.

i. **[3 points]** If both tables in a simple nested loop join fit entirely in memory, you must choose the larger relation as the inner relation to minimize I/O cost.
□ True □ False

ii. **[3 points]** If neither table fits entirely in memory, I/O costs would be lower if we process both tables on a per-block basis rather than per-tuple basis.
□ True □ False

iii. **[4 points]** A hash join is faster than a sort-merge join in all scenarios.
□ True □ False

iv. **[4 points]** An index nested loop join requires an index on the outer- and inner- tables.
□ True □ False

v. **[4 points]** For a hash join to work, the inner table (or its partitions) need to fit into memory.
□ True □ False

Question 3: Query Execution, Planning, and Optimization . . . . . [24 points]

(a) **[2 points]** The iterator model always allows tuples to fully pass through the entire sequence of operators before retrieving the next tuple.
☐ True ☐ False

(b) **[2 points]** Assume that the DBMS zone maps are up to date. The DBMS can use these zone maps to answer specific queries without reading any actual table heap tuples:
☐ True ☐ False

(c) **[2 points]** Assuming a query with multiple `OR` predicates. Using a multi-index scan will always perform better than a sequential scan.
☐ True ☐ False

(d) **[2 points]** For OLAP queries, which often operate on large datasets, it isn't worth parallelizing an individual query plan.
☐ True ☐ False

(e) **[2 points]** The process per DBMS worker approach provides better fault isolation and scheduling control than the thread per DBMS worker approach.
☐ True ☐ False

(f) **[2 points]** In OLAP workload, the vectorized model's performance improvements come mainly from the reduction in the number of disk I/O operations.
☐ True ☐ False

(g) **[2 points]** The query optimizer in a database management system always guarantees the generation of an optimal execution plan by exhaustively evaluating all possible plans to ensure the lowest cost for query execution.
☐ True ☐ False

(h) **[2 points]** Predicate and projection pushdown will always improve query performance.
☐ True ☐ False

(i) **[2 points]** The execution plan with the lowest cost is guaranteed to be the most efficient among all execution plans enumerated by the query optimizer.
☐ True ☐ False

(j) **[2 points]** Sampling statistics requires evaluating each tuple in the entire table.
☐ True ☐ False

(k) **[2 points]** Equi-depth histogram maintains counts for a group of values instead of each unique key to reduce memory footprint and uses the same range size for each bucket.
☐ True ☐ False

(l) A database contains a single table: `University(id,name,state,city)`. You need to estimate the cardinality of the following query:

```
SELECT * FROM University WHERE state = 'PA' AND city = 'Pittsburgh'
```

For the following questions, assume we know that `University` has 10,000 rows with 5% having `state = 'PA'`, and 0.5% having `city = 'Pittsburgh'`.

i. **[1 point]** Under uniform data assumption and independent predicates assumption, what is the estimated cardinality c of this query? Take $\lceil c \rceil$ of the result.
☐ 1 ☐ 2 ☐ 3 ☐ 10 ☐ 30

ii. **[1 point]** Is the result from previous question an overestimate or underestimate of the true cardinality?
☐ Overestimate ☐ Underestimate