

CARNEGIE MELLON UNIVERSITY
COMPUTER SCIENCE DEPARTMENT
15-445/645 – DATABASE SYSTEMS (FALL 2023)
PROF. ANDY PAVLO AND JIGNESH PATEL

Homework #3 (by Anurag Choudhary) – Solutions
Due: **Sunday, Oct 08, 2023 @ 11:59pm**

IMPORTANT:

- Enter all of your answers into **Gradescope by 11:59pm on Sunday, Oct 08, 2023.**
- **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 1 - 2 hours (0.5 - 1 hours for each question)

Revision : 2023/10/07 23:00

Question	Points	Score
Sorting Algorithms	36	
Join Algorithms	43	
Bloom Filters	21	
Total:	100	

Question 1: Sorting Algorithms [36 points]**Graded by:**

We have a database file with nine million pages ($N = 9,000,000$ 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) [6 points] Assume that the DBMS has fifty buffers. How many sorted runs are generated? Note that the final sorted file does not count towards the sorted run count.
 183,673 183,674 183,750 **183,751** 187,501 187,502

Solution:

$$\left\lceil \frac{9,000,000}{50} \right\rceil + \left\lceil \frac{180,000}{49} \right\rceil + \left\lceil \frac{3,674}{49} \right\rceil + \left\lceil \frac{75}{49} \right\rceil = 183,751$$

- (b) [6 points] Again, assuming that the DBMS has fifty buffers. How many pages (in general) does each sorted run have after the third pass (i.e. Note: this is Pass #2 if you start counting from Pass #0)?
 49 50 2,450 2,451 2,452 **120,050** 120,051
 120,052

Solution: On the first pass, B buffer pages will be used to create the sorted runs. From the second pass onward, $B-1$ runs will be sorted through a K -way merge.

First pass: 50 pages for each sorted run. Second pass: $50 * 49$ pages for each sorted run. Third pass: $50 * 49 * 49 = 120,050$ pages for each sorted run.

- (c) [6 points] Again, assuming that the DBMS has fifty buffers. How many passes does the DBMS need to perform in order to sort the file?
 1 2 3 4 **5**

Solution:

$$1 + \left\lceil \log_{B-1} \left(\left\lceil \frac{N}{B} \right\rceil \right) \right\rceil = 1 + \left\lceil \log_{49} (\lceil 9,000,000/50 \rceil) \right\rceil \\ = 1 + 4 = 5$$

- (d) [6 points] Suppose the DBMS has 100 buffers. What is the total I/O cost to sort the file?
 18,000,000 36,000,000 54,000,000 **72,000,000** 90,000,000

Solution: $\#Passes = 1 + \lceil \log_{99} (\lceil 9,000,000/100 \rceil) \rceil = 1 + 3 = 4$
 $Cost = 2N \times \#Passes = 2 \times 9,000,000 \times 4 = 72,000,000$

- (e) [6 points] What is the smallest number of buffers B such that the DBMS can sort the target file using only three passes?
- 55 56 208 209 3,000 3,001

Solution: We want the smallest integer B such that $N \leq B \times (B - 1)^2$. If $B = 209$, then $9,000,000 \leq 209 \times 208^2 = 9,042,176$; any smaller value for B would fail.

- (f) [6 points] Suppose the DBMS has twenty-three buffers. What is the largest database file (expressed in terms of the number of pages) that can be sorted with external merge sort using four passes?
- 10,648 11,132 12,167 234,256 244,904 279,841
 5,153,632 5,387,888 6,436,343 118,533,536

Solution: We want the largest integer N such that $N \leq B \times (B - 1)^3$. The largest such value is $B \times (B - 1)^3$ itself, which is $23 \times 22^3 = 244,904$

Question 2: Join Algorithms [43 points]**Graded by:**

Consider relations $X(a, c)$, $Y(a, b, e)$, 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 $B = 750$ pages in the buffer
- Table X spans $M = 2,500$ pages with 80 tuples per page
- Table Y spans $N = 600$ pages with 300 tuples per page
- Table Z spans $O = 1,500$ pages with 150 tuples per page
- The join result of X and Y spans $P = 500$ 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) [3 points] What is the I/O cost of a simple nested loop join with X as the outer relation and Y as the inner relation?

- 50,500 200,600 202,500 1,502,500 120,000,600
 120,002,500 300,002,500

Solution: $M + m \times N = 2500 + 2500 \times 80 \times 600 = 120,002,500$

- (b) [3 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?

- 1,500 2,100 2,700 3,000 **3,300** 3,600 3,900
 5,100

Solution: $O + \lceil \frac{O}{B-2} \rceil \times N = 1,500 + \lceil \frac{1,500}{748} \rceil \times 600 = 1,500 + 1,800 = 3,300$

- (c) [3 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?

- 1,500 **2,100** 2,700 3,000 3,300 3,600 3,900
 5,100

Solution: $N + \lceil \frac{N}{B-2} \rceil \times O = 600 + \lceil \frac{600}{748} \rceil \times 1,500 = 600 + 1,500 = 2,100$

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

- i. [3 points] What is the cost of sorting the tuples in X on attribute a ?

- 2,000 4,000 6,000 8,000 **10,000** 12,000

Solution: $passes = 1 + \lceil \log_{B-1}(\lceil \frac{M}{B} \rceil) \rceil = 1 + \lceil \log_{749}(\lceil \frac{2500}{750} \rceil) \rceil = 1 + 1 = 2$
 $2M \times passes = 2 * 2500 * 2 = 10,000$

- ii. **[3 points]** What is the cost of sorting the tuples in Z on attribute a?
 2,000 4,000 **6,000** 8,000 10,000 12,000

Solution: $passes = 1 + \lceil \log_{B-1}(\lceil \frac{O}{B} \rceil) \rceil = 1 + \lceil \log_{749}(\lceil \frac{1,500}{750} \rceil) \rceil = 1 + 1 = 2$
 $2O \times passes = 2 * 1,500 * 2 = 6,000$

- iii. **[3 points]** What is the cost of the merge phase in the worst-case scenario?
 1,500 2,000 2,500 4,000 900,000 2,000,000
 3,250,000 **3,750,000** 4,500,000 5,000,000

Solution: $M \times O = 2,500 \times 1,500 = 3,750,000$

- iv. **[3 points]** What is the cost of the merge phase assuming there are no duplicates in the join attribute?
 1,500 2,000 2,500 **4,000** 900,000 2,000,000
 3,250,000 3,750,000 4,500,000 5,000,000

Solution: $M + O = 2,500 + 1,500 = 4,000$

- v. **[3 points]** Now consider joining X, Y and then joining the result with Z. What is the cost of the final merge phase assuming there are no duplicates in the join attribute?
 1,000 **2,000** 3,000 5,000 1,000,000

Solution: $P + O = 500 + 1,500 = 2,000$

- (e) Consider a hash join with Y as the outer relation and Z as the inner relation. You may ignore recursive partitioning and partially filled blocks.

- i. **[3 points]** What is the cost of the probe phase?
 1,000 2,000 **2,100** 4,200 6,400 7,200 10,000

Solution: $(N + O) = (600 + 1,500) = 2,100$

- ii. **[3 points]** What is the cost of the partition phase?
 1,000 2,000 2,100 **4,200** 6,400 7,200 10,000

Solution: $2 \times (N + O) = 2 \times (600 + 1,500) = 2 \times 2,100 = 4,200$

- (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 h_1 for large buckets.

Use linear probing for collisions and page in and out parts of the hashtable needed at a given time.

■ **Create hashables for the inner and outer relation using h_1 and rehash into an embedded hash table using $h_2 \neq h_1$ for large buckets.**

Create two hashables half the size of the original one, run the same hash join algorithm on the tables, and then merge the hashables together.

Solution: Use Grace hash join with recursive partitioning, which is what the correct option describes.

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

i. [2 points] In a simple nested loop join where both tables fit entirely in memory, the choice of which table to use as the inner/outer table significantly affects I/O costs.

True ■ **False**

Solution: If both tables fit entirely in memory, then they can be read just once and therefore the order would not be very important.

ii. [2 points] In a simple nested loop join where one of the tables fits entirely in memory, it is beneficial to use that table as the inner table.

■ **True** False

Solution: Using the table that fits in memory as the inner table would reduce I/O costs as it would then be read only once, instead of repeatedly for each tuple of the outer table.

iii. [2 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

Solution: A block nested loop join has fewer disk accesses when compared to a simple nested loop join.

iv. [2 points] For a block nested loop join, in the worst case, each block in the inner table has to be read once for each tuple in the outer table.

True ■ **False**

Solution: Even in the worst case, each block in the inner table is read only once for each *block* in the outer table.

v. [2 points] In an index nested loop join where only one of the tables has an index on the join attribute, the choice of inner/outer tables should be made based on the table

sizes.

True **False**

Solution: The choice of the inner/outer tables would be affected by which table has the index on the join attribute. If both tables had indices, then their sizes would have been a factor.

Question 3: Bloom Filters.....[21 points]**Graded by:**

Assume that we have a bloom filter that is used to register names. The filter uses two hash functions h_1 and h_2 which hash the following strings to the following values:

input	h_1	h_2
“Nas”	1973	1994
“Cole”	1985	2014
“Kendrick”	1987	2015
“Rocky”	1988	2013
“JID”	1990	2022
“Denzel”	1995	2018

(a) [4 points] Suppose the filter has 6 bits initially set to 0:

bit 0	bit 1	bit 2	bit 3	bit 4	bit 5
0	0	0	0	0	0

Which bits will be set to 1 after “Cole” and “JID” have been inserted?

0 1 2 3 4 5

Solution: Because the filter has 6 bits, we take the modulo of the hashed output and 6.

“Cole”:

$$1985 \bmod 6 = 5$$

$$2014 \bmod 6 = 4$$

“JID”:

$$1990 \bmod 6 = 4$$

$$2022 \bmod 6 = 0$$

(b) [3 points] If “Denzel” is inserted next, which bits will now be set to 1 (including those already set in part (a))?

0 1 2 3 4 5

Solution: $1995 \bmod 6 = 3$

$$2018 \bmod 6 = 2$$

Bits 0, 2, 3, 4, 5 are now set.

(c) [3 points] What will the filter return if we now lookup “Nas”?

True False

Solution: $1973 \bmod 6 = 5$

$$1994 \bmod 6 = 2$$

Because bit 2 and bit 5 are both 1, the filter returns true.

(d) [3 points] Suppose the filter has 6 bits set to the following values:

bit 0	bit 1	bit 2	bit 3	bit 4	bit 5
0	1	0	1	0	1

What will the filter return if we lookup “Rocky”?

True **False**

Solution: $1988 \bmod 6 = 2$

$2013 \bmod 6 = 3$

Bits 2 and 3 must both be 1 if “Rocky” has been inserted. Because bit 2 is 0, the filter returns false.

(e) [3 points] What will the filter from part (d) return if we lookup “Kendrick”?

True False

Solution: $1987 \bmod 6 = 1$

$2015 \bmod 6 = 5$

Because bit 1 and bit 5 are both 1, the filter returns true.

(f) [5 points] Suppose the filter has 6 bits set to the following values:

bit 0	bit 1	bit 2	bit 3	bit 4	bit 5
0	0	1	1	0	1

Which names are guaranteed to have NOT been inserted?

“Nas” “Cole” “Kendrick” “Rocky” “JID” “Denzel”

Solution: The bits that must be set for each name are:

input	first bit	second bit	inserted?
“Nas”	5	2	Maybe
“Cole”	5	4	No
“Kendrick”	1	5	No
“Rocky”	2	3	Maybe
“JID”	4	0	No
“Denzel”	3	2	Maybe

“Nas” might have been inserted since both bit 2 and bit 5 are set (in this case - definitely been inserted).

One or both of “Rocky” and “Denzel” have been inserted. We cannot guarantee that either has not been inserted.