

University of Waterloo

Midterm Examination

Winter, 2012

Student Name: _____

Student ID Number: _____

Course Abbreviation and Number	CS448
Course Title	Database Systems Implementation
Sections	01 (13:00)
Held With Courses	CS 648
Sections of Held With Courses	01
Instructor	K. Salem

Date of Exam	February 16, 2011
Time Period	19:00-21:00
Duration of Exam	120 minutes
Number of Exam Pages (including this cover sheet)	8 pages
Exam Type	Closed Book
Additional Materials Allowed	None

Question 1: (4 marks)	Question 2: (8 marks)	Question 3: (8 marks)	Question 4: (9 marks)
Question 5: (10 marks)	Question 6: (8 marks)	Question 7: (3 marks)	
Total: (50 marks)			

1. (4 total marks)

Suppose that a relational database contains the following large relation:

Track(ReleaseID, TrackNum, Title, Length, Location)

The database has a clustered B+-tree index on the composite key **ReleaseID,TrackNum**. It also has an unclustered B+-Tree index on **TrackNum**.

- a. (2 marks)** Give an example of a simple, single-table SQL query (not an update, insertion, or deletion) that would require fewer I/O operations to answer using the clustered index than the unclustered index. Your example should also require fewer I/O operations to answer using the clustered index than using a full table scan, i.e., than not using either index. For full credit, explain briefly (in one or two sentences) why the query would require less I/O if the clustered index were used.
- b. (2 marks)** Give an example of a simple, single-table SQL query (not an update, insertion, or deletion) that would require fewer I/O operations to answer using the unclustered index than the clustered index. Your example should also require fewer I/O operations to answer using the clustered index than using a full table scan, i.e., than not using either index. For full credit, explain briefly (in one or two sentences) why the query would require less I/O if the unclustered index were used.

2. (8 total marks)

Suppose that a database system has a hash index on a relation R . The index is implemented using extensible hashing. Suppose that a client inserts a new record into R , and that the insertion causes the size of the hash index directory to double, from n entries to $2n$ entries. For the purposes of this question, assume that there have not been any deletions from this index - only insertions.

a. (2 marks)

Immediately prior to the insertion that doubles the directory size, how many hash buckets did this index have? (Hash buckets are what the directory entries point to.) Give two answers: the maximum number of hash buckets it might have had, and the minimum number of hash buckets it might have had. Express your answers in terms of n .

b. (2 marks)

Let k represent the exact number of hash buckets in the index immediately before the insertion that doubles the directory size. How many hash buckets will the index have immediately after the insertion that doubles the directory size? Express your answer in terms of k .

c. (2 marks)

Immediately after the insertion that doubles the directory size, how many hash buckets will have exactly one directory entry pointing to them?

d. (2 marks) Immediately after the insertion that doubles the directory size, what is the maximum number of directory entries that could be pointing at a single hash bucket in the index?

3 (8 total marks)

Consider a database that contains the following two relations

```
Track(TrackID,Title,ReleaseID,Duration,Format,Genre)
Release(ReleaseID,Title,ArtistID,ReleaseDate)
```

The DBMS executes the following query:

```
select T.Title,T.Genre,R.ReleaseDate,R.ReleaseID
from Track T, Release R
where T.ReleaseID = R.ReleaseID
and T.Format = 'mp3'
```

a. (4 marks)

Suppose that the DBMS executes this query using a block nested loop join iterator with **Track** as the outer relation and **Release** as the inner. Furthermore, suppose that the access method for **Track** produces **Track** tuples in **Genre** order, and the access method for **Release** produces **Release** tuples in **ReleaseDate** order. Which of the following statements must be true? (Identify all that are true.)

- The output of the join will be in **ReleaseDate** order.
- The output of the join will be in **Genre** order.
- The output of the join will be in **ReleaseID** order.
- The join will consume all **Track** tuples before producing any output tuples.
- The join will consume all **Release** tuples before producing any output tuples.
- The join will consume all **Track** tuples before consuming any **Release** tuples.
- The join will consume all **Release** tuples before consuming any **Track** tuples.

b. (4 marks)

Suppose instead that the DBMS executes this query using a hybrid hash join iterator with **Track** as the probe relation and **Release** as the build relation. As was the case for part (a), suppose that the access method for **Track** produces **Track** tuples in **Genre** order, and the access method for **Release** produces **Release** tuples in **ReleaseDate** order. Which of the following statements must be true? (Identify all that are true.)

- The output of the join will be in **ReleaseDate** order.
- The output of the join will be in **Genre** order.
- The output of the join will be in **ReleaseID** order.
- The join will consume all **Track** tuples before producing any output tuples.
- The join will consume all **Release** tuples before producing any output tuples.
- The join will consume all **Track** tuples before consuming any **Release** tuples.
- The join will consume all **Release** tuples before consuming any **Track** tuples.

- 4 (9 total marks)** Suppose that a database contains two relations, R and S . R has 50,000 tuples and occupies 500 data blocks, and has a B+-tree index on attribute $R.a$. The index is unclustered and has a height of 2. The leaves of the index occupy 10 blocks. The relation S has 200,000 tuples and occupies 2000 data blocks.

Suppose that the DBMS executes the following three queries in sequence, i.e., one at a time, in the order listed:

Q_1 : `select * from R order by R.a`

Q_1 is executed by using the index to read the tuples of R in $R.a$ order.

Q_2 : `select * from S`

Q_2 is executed by using a table scan of S .

Q_3 : `select * from R order by R.a`

Q_3 is the same as Q_1 , and is executed the same way.

a. (3 marks)

Assuming that the buffer cache is very small (no cache hits), how many disk block reads will the DBMS require to execute each of the three queries? Your answer should consist of three numbers, one for each query.

b. (3 marks)

Repeat part (a), but this time assume that the buffer cache is large enough to hold 1200 blocks, and that the buffer replacement policy is LRU. Assume that the buffer cache is empty prior to the execution of Q_1 . Again, be sure to give an answer for each query.

c. (3 marks)

Repeat part (a), but this time assume that the buffer cache is large enough to hold 1200 blocks, and that the buffer replacement policy is 2Q. (2Q uses a target size of 300 for its FIFO queue.) Assume that the buffer cache is empty prior to the execution of Q_1 . Be sure to give an answer for each query.

- 5 (10 total marks)** Suppose that a database contains a relation R stored in B blocks, with T tuples per block. The DBMS executes the following query:

```
select * from R where R.a > 10 order by R.a
```

Let f represent the fraction of R 's tuples that satisfy the query predicate $R.a > 10$. For example, $f = 0$ means that none of R 's tuples satisfy the predicate, $f = 0.5$ means that half of R 's tuples satisfy the predicate, and $f = 1.0$ means that all of the tuples satisfy the predicate.

Suppose that the DBMS executes the query using the following plan:

1. use a table scan to read all of R , filtering out on the fly the non-matching tuples (those that do not satisfy the predicate).
2. sort the matching tuples, using an external merge sort, to produce them in $R.a$ order.

a. (4 marks)

What is the I/O cost (number of disk block I/O operations) required to execute this plan? Indicate the cost of each iterator (table scan and sort) separately. Assume that the merge sort operator has sufficient memory to complete the sort using a single merging phase. Express your answer in terms of f and the other constants listed above.

b. (3 marks)

Let M represent the amount of memory (measured in blocks) available to the merge sort operator. What is the minimum value of M required to ensure that the merge sort can complete the sort using a single merging pass? Express your answer in terms of f and the other constants listed above.

c. (3 marks)

If R has an unclustered index on $R.a$, an alternative plan for executing the query would be to retrieve matching tuples in $R.a$ order using the index, thus eliminating the sort operation. As f approaches 0.0, which of the two plans (table scan then sort, or index scan) do you expect will have the lowest cost? Briefly explain your answer.

6 (8 total marks) Provide brief answers to the each of the following questions. “Brief” means a sentence or two.

a. (2 marks)

Why does a DBMS buffer manager pin pages, and what effect does pinning have on the buffer’s page replacement policy?

b. (2 marks)

Many DBMS store data in files. Identify at least two ways that an underlying file system can impact the performance or correctness of the DBMS.

c. (2 marks)

Suppose that an iterator-based query plan includes hash join, and that the join has enough memory so that it only needs to partition its inputs once, i.e, there is no recursive partitioning. How many times will each tuple of the build input be written to and read from secondary storage by the hash join operator? How many times will each tuple of the probe input be written to and read from secondary storage by the hash join operator?

d. (2 marks)

In pipelined query plans, some operators (like external merge sort) are called blocking operators. What does it mean for an operator to be a blocking operator?

- 7 (3 marks)** Suppose that a database occupies a total of P pages (blocks) on secondary storage. The database is stored on 10 disk drives, each of which cost \$100 and provides 100 IOPS (I/O Operations per Second). Suppose that, given the current price of memory and the size of the pages, an analysis like the one that led to the 5 Minute Rule determines that M ($0 < M < P$) of the database pages should be kept in memory, and the remaining $P - M$ pages should be kept on disk.

Now, suppose that the 10 disk drives that hold the database are replaced with a single SSD (flash memory based) storage device. This new device costs \$500 and provides 50,000 IOPS. (It is large enough to hold the entire database if necessary.) The 5 Minute Rule analysis is repeated for the new system, and the analysis determines that M_{SSD} ($0 < M_{SSD} < P$) of the database pages should be kept in memory. Assuming that the cost of memory and the database pages' access patterns have not changed, which one of the following must be true?

- $M_{SSD} = M$
- $M_{SSD} \geq M$
- $M_{SSD} \leq M$

Briefly justify your answer.