
Database Systems

Instructors: Hao-Hua Chu
Winston Hsu
Fall Semester, 2007

Assignment 6

Deadline: At the end of the class, Jan 7 (Mon), 2008

This is an individual assignment, that is, no group submissions are allowed.

Cheating Policy: If you are caught cheating, your grade is 0.

Late Policy: We will not accept any assignment submissions.

Questions

- I. Consider a disk with an average seek time of 10ms, average rotational delay of 5ms, and a transfer time of 1ms for a 4K page. Assume that the cost of reading/writing a page is the sum of these values (i.e., 16ms) unless a sequence of pages is read/written. In this case, the cost is the average seek time plus the average rotational delay (to find the first page in the sequence) plus 1ms per page (to transfer data). You are given 320 buffer pages and asked to sort a file with 10,000,000 pages.
 1. Why is it a bad idea to use the 320 pages to support virtual memory, that is, to *new* 10,000,000 * 4K bytes of memory, and to use an in-memory sorting algorithm such as Quicksort?
 2. Assume that you begin by creating sorted runs of 320 pages each in the first pass. Evaluate the cost of the following approaches for the subsequent merging passes:
 - (a) Do 319-way merges.
 - (b) Create 256 'input' buffers of 1 page each, create an 'output' buffer of 64 pages, and do 256-way merges.
 - (c) Create 16 'input' buffers of 16 pages each, create an 'output' buffer of 64 pages, and do 16-way merges.
 - (d) Create eight 'input' buffers of 32 pages each, create an 'output' buffer of 64 pages, and do eight-way merges.
 - (e) Create four 'input' buffers of 64 pages each, create an 'output' buffer of 64 pages, and do four-way merges.

- II. Consider the join $R \bowtie_{R.a=S.b} S$, given the following information about the relations to be joined. The cost metric is the number of page I/Os unless otherwise noted, and the cost of writing out the result should be uniformly ignored.

Relation R contains 10,000 tuples and has 10 tuples per page.

Relation S contains 2000 tuples and also has 10 tuples per page.

Attribute b of relation S is the primary key for S.

Both relations are stored as simple heap files.

Neither relation has any indexes built on it.

52 buffer pages are available.

1. What is the cost of joining R and S using a page-oriented simple nested loops join? What is the minimum number of buffer pages required for this cost to remain unchanged?
 2. What is the cost of joining R and S using a block nested loops join? What is the minimum number of buffer pages required for this cost to remain unchanged?
 3. What is the cost of joining R and S using a sort-merge join? What is the minimum number of buffer pages required for this cost to remain unchanged?
 4. What is the cost of joining R and S using a hash join? What is the minimum number of buffer pages required for this cost to remain unchanged?
 5. What would be the lowest possible I/O cost for joining R and S using any join algorithm, and how much buffer space would be needed to achieve this cost? Explain briefly.
 6. How many tuples does the join of R and S produce, at most, and how many pages are required to store the result of the join back on disk?
 7. Would your answers to any of the previous questions in this exercise change if you were told that R.a is a foreign key that refers to S.b?
- III. Consider a database with objects X and Y and assume that there are two transactions T1 and T2. Transaction T1 reads objects X and Y and then writes object X. Transaction T2 reads objects X and Y and then writes objects X and Y.
1. Give an example schedule with actions of transactions T1 and T2 on objects X and Y that results in a write-read conflict.
 2. Give an example schedule with actions of transactions T1 and T2 on objects X and Y that results in a read-write conflict.
 3. Give an example schedule with actions of transactions T1 and T2 on objects X and Y that results in a write-write conflict.
 4. For each of the three schedules, show that Strict 2PL disallows the schedule.

- IV. Consider the following classes of schedules: *serializable*, *conflict-serializable*, *view-serializable*, recoverable. For each of the following schedules, state which of the preceding classes it belongs to. If you cannot decide whether a schedule belongs in a certain class based on the listed actions, explain briefly.

The actions are listed in the order they are scheduled and prefixed with the transaction name. If a commit or abort is not shown, the schedule is incomplete; assume that abort or commit must follow all the listed actions.

1. T1:R(X), T2:R(X), T1:W(X), T2:W(X)
2. T1:W(X), T2:R(Y), T1:R(Y), T2:R(X)
3. T1:R(X), T2:R(Y), T3:W(X), T2:R(X), T1:R(Y)
4. T1:R(X), T1:R(Y), T1:W(X), T2:R(Y), T3:W(Y), T1:W(X), T2:R(Y)
5. T1:R(X), T2:W(X), T1:W(X), T2:Abort, T1:Commit
6. T1:R(X), T2:W(X), T1:W(X), T2:Commit, T1:Commit
7. T1:W(X), T2:R(X), T1:W(X), T2:Abort, T1:Commit
8. T1:W(X), T2:R(X), T1:W(X), T2:Commit, T1:Commit
9. T1:W(X), T2:R(X), T1:W(X), T2:Commit, T1:Abort
10. T2: R(X), T3:W(X), T3:Commit, T1:W(Y), T1:Commit, T2:R(Y), T2:W(Z), T2:Commit
11. T1:R(X), T2:W(X), T2:Commit, T1:W(X), T1:Commit, T3:R(X), T3:Commit
12. T1:R(X), T2:W(X), T1:W(X), T3:R(X), T1:Commit, T2:Commit, T3:Commit

Submission

Hand in PAPER PRINTOUT that contains your answers to the four questions. Please include your name and ID.