

CS W186 Databases - Fall 2019
Guerilla Section 2: Sorting, Hashing, and Buffer Management

Sunday, September 29, 2019

1 Sorting and Hashing

Suppose the size of a page is 4KB, and the size of the memory buffer is 1 MB (1024 KB).

1. We have a relation of size 800 KB. How many page IOs are required to sort this relation?
2. We have a relation of size 5000 KB. How many page IOs are required to sort this relation?
3. What is the size of the largest relation that would need two passes to sort?
4. What is the size of the largest relation we can possibly hash in two passes (i.e. with just one partitioning phase)?

5. Suppose we have a relation of size 3000 KB. We are executing a `DISTINCT` query on a column `age`, which has only two distinct values, evenly distributed. Would sorting or hashing be better here, and why?
6. Now suppose we were executing a `GROUP BY` instead. Would sorting or hashing be better here, and why?

2 External Sorting

Assume our buffer pool has 8 frames. In this question, we'll externally sort a 500 page file.

1. How many passes will it take to sort this file?
2. Given the number of passes you calculated in 2.1, how many I/Os are necessary to externally sort the file?
3. What is the minimum number of additional frames needed to reduce the number of passes found in 2.1 by 1?
4. What is the minimum number of additional frames needed to sort the file in one pass?

3 Buffer Management

We're given a buffer pool with 4 pages, which is empty to begin with. Answer the following questions given this access pattern:

A B C D E B A D C A E C

1. What is the hit rate for MRU?
2. In order of when they were first placed into the buffer pool, what pages remain in the buffer pool after this sequence of accesses?
3. What is the hit rate for clock?
4. Which pages are in the buffer pool after this sequences of accesses?
5. Which pages have their reference bits set?
6. Which page is the arm of the clock pointing to?

4 Hints for Question 1: Hashing and Sorting

1. Convert all numbers to pages, and then apply the equation you learned in class.
2. Same.
3. Again, convert all numbers to pages.
The second pass merges how many runs?
For the second pass to complete the sort, this must be enough to merge all the runs created by the first sort. How big does the relation have to be for the first sort to create that many runs?
4. Similar to the previous question, but think about the hashing protocol instead.
5. Which protocol benefits from removing duplicates?
6. `GROUP BY` does not allow us to remove duplicates. Are both protocols capable of handling such large numbers of duplicates?