


Tutorial on Sorting

Database System Technology

Niv Dayan



Midterm covers all
material so far
(Including this week)




Open book

Office Hours


TA Office Hours
4-6 pm
(Project + content)

Instructor Office Hours
4-5 pm and 7-8 pm


Suppose we need next min entry from buffer 2 but it is empty.




Suppose we need next min entry from buffer 2 but it is empty.





Double buffering: load one additional buffer preemptively for each partition before the first buffer empties



Double buffering: load one additional buffer preemptively for each partition before the first buffer empties



Larger though fewer buffers: more groups, so potentially more iterations,
but each I/O reads more data



Question 1

Suppose you have a file with 10,000 pages and you have 17 buffers in memory, each of them 4KB. We need to externally sort the file.



17 buffers



Table with 10,000 pages


Question 1

Suppose you have a file with 10,000 pages and you have 17 buffers in memory, each of them 4KB. We need to externally sort the file.

- (A) How many partitions are created after the first pass?
- (B) How many passes does it take to sort the file completely?
- (C) How many I/Os are issued in total to sort the file?
- (D) How many buffers would you need to sort the file in just 2 passes?

Question 2

Consider a table that is too large to fit in memory. We have an unclustered B-tree index over column A. We get the following query: “Select * from table order by A”. There are two options to process this query: (1) scan the index, or (2) externally sort the file based on column A. What are the costs of these methods, and under which circumstances would you choose each one?



Select * from table sort by A

Question 2

Consider a table that is too large to fit in memory. We have an unclustered B-tree index over column A. We get the following query: “Select * from table order by A”. There are two options to process this query: (1) scan the index, or (2) externally sort the file based on column A. What are the costs of these methods, and under which circumstances would you choose each one?

Follow up: how would your answer change if we have a clustered index on column A?

Question 3

You are given M memory and N entries where $N \gg M \gg 3B$. Why is it a bad idea to use the available memory as virtual memory (e.g., to allocate using 'new' space for M entries, and to use in-memory Quicksort)? Use cost models to justify your answer.

Question 4

Suppose we have a tiered LSM-tree with a size ratio of T . Analyze the CPU costs of compaction assuming we check all buffers each time to find the minimum key. Then propose a technique to improve CPU costs.

