

CS43002: Database Management Systems: Class Test II

Name:

Roll:

1. Consider a disk system with 100 cylinders. The requests to access the cylinders occur in following sequence:

4, 34, 10, 7, 19, 73, 2, 15, 6, 20

143

Assuming that the head is currently at cylinder 1, what is the time taken to satisfy all requests? It takes 1ms to move from one cylinder to adjacent one and *elevator scheduling* policy is used?

2. A file is organized so that the ordering of data records is the same as or close to the ordering of data entries in some index. Then that index is called

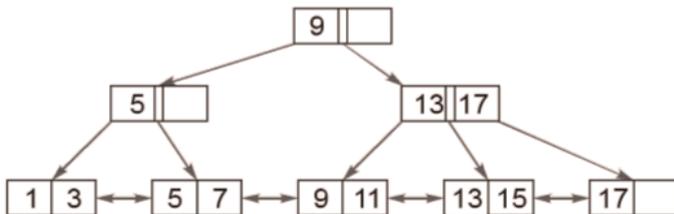
Clustered/Primary

3. A B+ -tree index is to be built on the *Name* attribute of the relation *STUDENT*. Assume that all student names are of length 8 bytes, disk block are size 512 bytes, and index pointers are of size 4 bytes. Given this scenario, what would be the best choice of the degree (i.e. the number of pointers per non-leaf node) of the B+ -tree?

43

4. With reference to the B+ tree index of order 1 shown below, the minimum number of nodes (including the root node) that must be fetched in order to satisfy the following query: "Get all records with a search key greater than or equal to 7 and less than 15" is?

5



5. Which one of the following hash functions on integers will distribute keys most uniformly over 10 buckets numbered 0 to 9 for i ranging from 0 to 2020?

B

(A) $h(i) = i^2 \bmod 10$ (B) $h(i) = i^3 \bmod 10$ (C) $h(i) = (11 * i^2) \bmod 10$ (D) $h(i) = (12 * i) \bmod 10$

6. Consider a join between relations r and s using the *nested loop* method. There are 3 buffers each of size equal to disk block size, out of which one buffer is reserved for intermediate results. Assuming $\text{size}(r) < \text{size}(s)$, which of the relation should be in outer loop so that the join will have fewer number of disk block accesses.

 r

7. Let relations $r1(A,B,C)$ and $r2(C,D,E)$ have the following properties: $r1$ has 20,000 tuples, $r2$ has 45,000 tuples, 25 tuples of $r1$ fit on one block, and 30 tuples of $r2$ fit on one block. Buffer size is 201 blocks. Estimate the number of disk accesses required by the *nested loop join* method, with $r1$ as the outer loop, for computing $r1 \bowtie r2$:

$$20000 * 1500 + 800 = 30,000,800$$

8. For the data of previous question (Q7) estimate the number of disk accesses required by the *block nested loop join* method, with $r1$ as the outer loop, for computing $r1 \bowtie r2$:

$$800/200 * 1500 + 800 = 6800$$

9. Consider a table r with attributes A_1, \dots, A_n . Let c_1 and c_2 be two constants. The following two queries can be evaluated either using "ordered indexing" or "hashing" on A_1 . Which one of the following statement is true?

Q1. $\Pi_{A_1 \dots A_n}(\sigma_{A_1=c_1}(r))$

Q2. $\Pi_{A_1 \dots A_n}(\sigma_{c_1 \leq A_1 \leq c_2}(r))$

C

- (A) Ordered indexing will always outperform hashing for both queries
- (B) Hashing will always outperform ordered indexing for both queries
- (C) Hashing will outperform ordered indexing on Q1, but not on Q2
- (D) Hashing will outperform ordered indexing on Q2, but not on Q1

10. The following functional dependencies hold for relations $R(A, B, C)$ and $S(B, D, E)$: $B \rightarrow A$ and $A \rightarrow C$.

Relations R contains 200 tuples and S contain 100 tuples. What is the maximum number of tuples possible in the natural join $R \bowtie S$.

100

----- END -----