

You will submit your solution to this assignment to the Curator System (as HW03). Your solution must be either a plain text file (e.g., NotePad) or a typed MS Word document; submissions in other formats will not be graded.

Except as noted, credit will only be given if you show relevant work.

1. [7 points each] Assume that we have a hash table of 10 slots, and we use the hash function $h(k) = k \bmod 10$. Show the results of inserting the following series of records in the table using the various collision resolution methods: 11, 31, 32, 56, 46, 47, 66.

- (a) Linear probing with step size 1

0	1	2	3	4	5	6	7	8	9

- (b) Linear probing with step size 3

0	1	2	3	4	5	6	7	8	9

- (c) Quadratic probing with k^2

0	1	2	3	4	5	6	7	8	9

- (d) Double hashing with second hash function $h_2(k) = k/10$.

0	1	2	3	4	5	6	7	8	9

2. [8 points each] The designer of a hash table decides to use some sort of probing strategy to resolve collisions. The exact probing strategy is not important to this question. The designer also decides to never allow the table to become complete full, by resizing/rehashing before that happens.

The designer decides to avoid the use of tombstones when deletions are performed. Instead, he proposes to use only two slot states, empty and full, and apply the following approach when a deletion is performed:

Suppose that the element to be deleted is in slot L ; use the probing strategy, taking one (more) probe step from slot L , to see if there is an element that was pushed beyond slot L (i.e., the one due to a collision with the now-deleted element in slot L). If so, move that element into slot L and delete it from its original slot, applying this same strategy from that position in the table. Stop when probing reaches a slot that is empty or that contains an element that did not collide with the element being deleted.

The designer argues that this will not only simplify the design of the table, but it will make searching more efficient, in some cases, since it will sometimes move elements "closer" to their home slot.

- Explain how the designer can determine whether the element that is "one more probe step from slot L " has been "pushed beyond slot L " due to a collision with the element that was formerly in slot L .
- Explain why the strategy will not yield a working hash table, at least in some situations. Efficiency is not the issue here; logical correctness is.

3. [7 points each] Assume a system uses a hard drive with the following physical characteristics:

total capacity	64 GB
# of platters	4
# of tracks per surface	16384
# of sectors per track	2048
cluster size	4 KB
spindle speed	14400 RPM
head start time	0.1 ms
track to track seek time	0.01 ms

In answering the following questions, express all final time values to the nearest thousandth of a millisecond (e.g, 8.333 ms). For full credit, you must show and explain all of your calculations.

- a) What is the average random head seek time for this drive?
 - b) What is the average rotational latency for this drive?
 - c) What is the average total time required to read one randomly-chosen sector from this drive?
 - d) What is the average total time required to read one randomly-chosen cluster from this drive?
 - e) What is the average total time required to read a file of 20 MB from this drive if the clusters are stored as efficiently as possible on the drive?
 - f) What is the average total time required to read a file of 20 MB from this drive if the clusters are randomly scattered on the drive?
4. [7 points each] Assume each of the changes described below were made to the design of the drive described in the preceding question, and that no other physical changes were made. Indicate the effect of the change on the average random head seek time and the average rotational latency time for the drive. The parts are independent. Justify your conclusions if you want credit.
- a) doubling the rotational speed, and modifying the read/write heads so that they can keep up with the increased transfer rate
 - b) reducing the platter diameter by half and doubling the number of tracks per surface; the track-to-track seek time is not changed