

1. BST, AVL trees, and hash tables can all be used to implement a dictionary ADT.

Dictionary Successful Search Comparisons with 10,000 integer items (Time in seconds)						
	Items added in sorted order		Items added in random order		Order did not matter (Hash table sizes $2^{15} = 32K$ )	
	BST	AVL Tree	BST	AVL Tree	Open Addr. (Quadratic)	Closed Addr. (Chaining)
Total add/put time	47.785	0.205	0.119	0.195	0.064	0.074
Total search time	38.100	0.060	0.079	0.062	0.044	0.039
Height of resulting tree	9,999	13	30	15	NA	NA


a) The puts of these 10,000 randomly ordered items into the BST took 0.119 seconds and 0.195 seconds into the AVL tree. Why did the BST puts take less time even though the final height was 30 vs. a final AVL tree height of 15?

*pivoting takes time in AVL tree  $O(\log_2 n)$   $O(1)$*


*$$\left[ \text{home add} + \frac{(\text{rehash}^2 + \text{rehash}^2)}{2} \right] \text{ hash table size}$$*

b) With a very, very poor hash function or very, very bad choice of keys all keys could hash to the same home address.

- What would be the worst-case big-oh of open-address hashing with quadratic probing?  $O(n)$
- What would be the worst-case big-oh of chaining using a linked list at each home address (i.e., ChainingDict)?  $O(n)$
- What would be the worst-case big-oh of chaining using an AVL tree at each home address?  $O(n)$



$O(\log_2 n)$



2. The data structures we have discussed so far are all in-memory, i.e., data is stored in main/RAM memory. Data can also be stored on secondary storage in a file (e.g., movieData.txt file). Currently, most secondary storage consists of hard-disks.

a) Complete the following table comparing main/RAM memory vs. ~~hard-disk~~ Secondary Storage

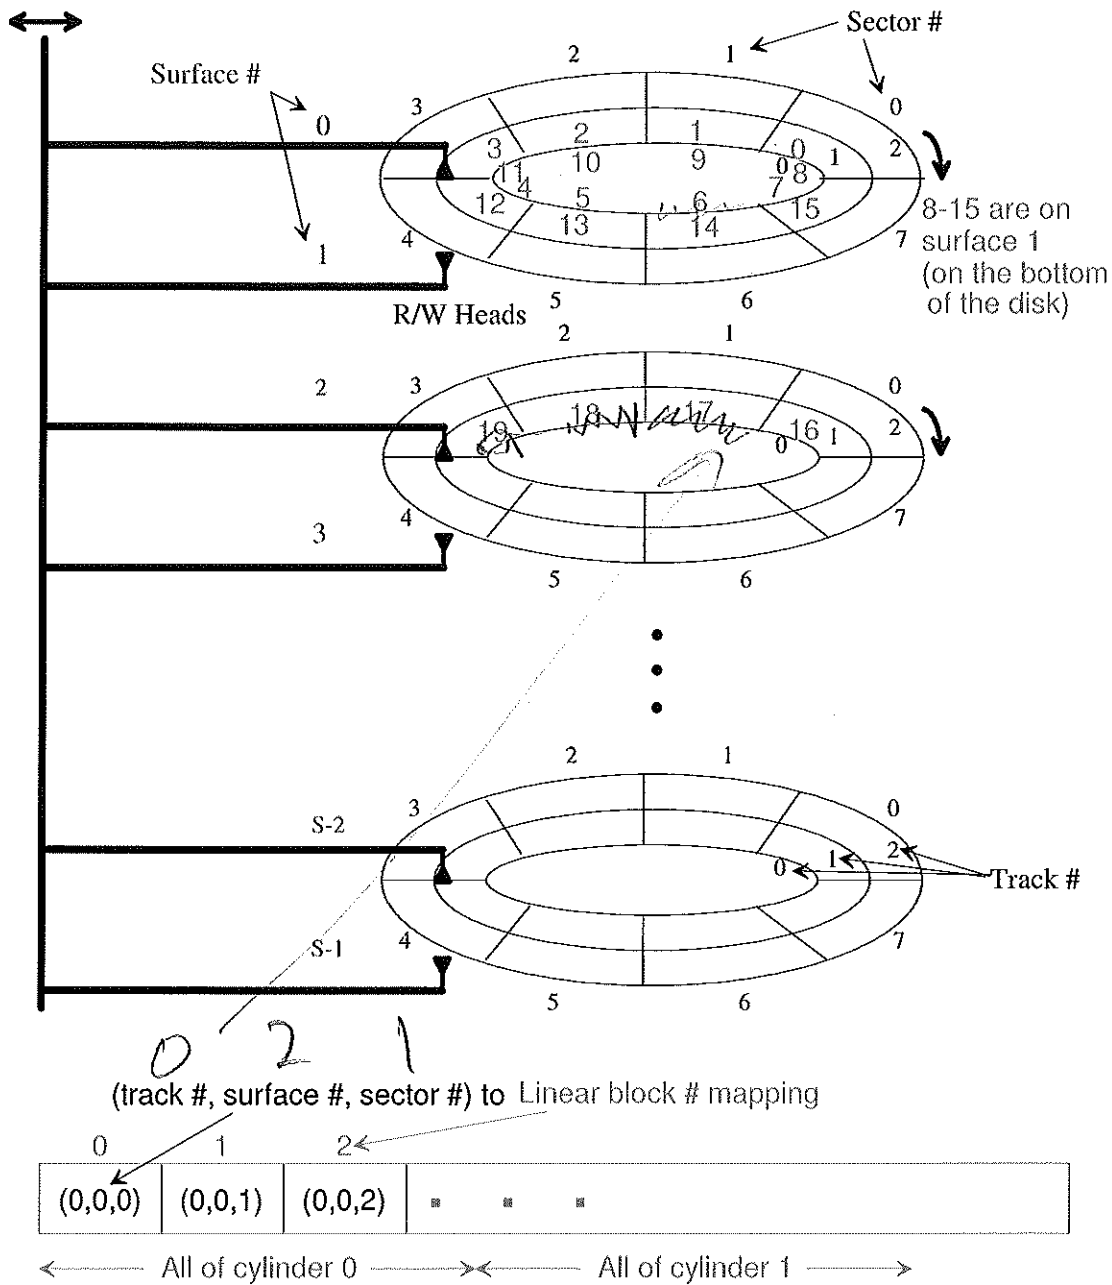
Criteria	Main/RAM memory	Hard-disk Drive	Solid-State Drive
Size on a typical desktop computer	4-8 GB	1 TB (1000GB)	250 GB
Average access time	10-50 nsec.	10ms	100 μsec?

b) Which criterion seems to be the most important difference between the main and secondary memories?  $n \cdot 10^{-9}$ ,  $\mu \cdot 10^{-6}$ ,  $m \cdot 10^{-3}$

*Size minor  
speed major  $\Rightarrow$  we don't want to*

*do many accesses to 2nd storage.* Lecture 24 Page 1

Logical View of Disk as Linear Collection of Blocks



Bits of linear block #: 

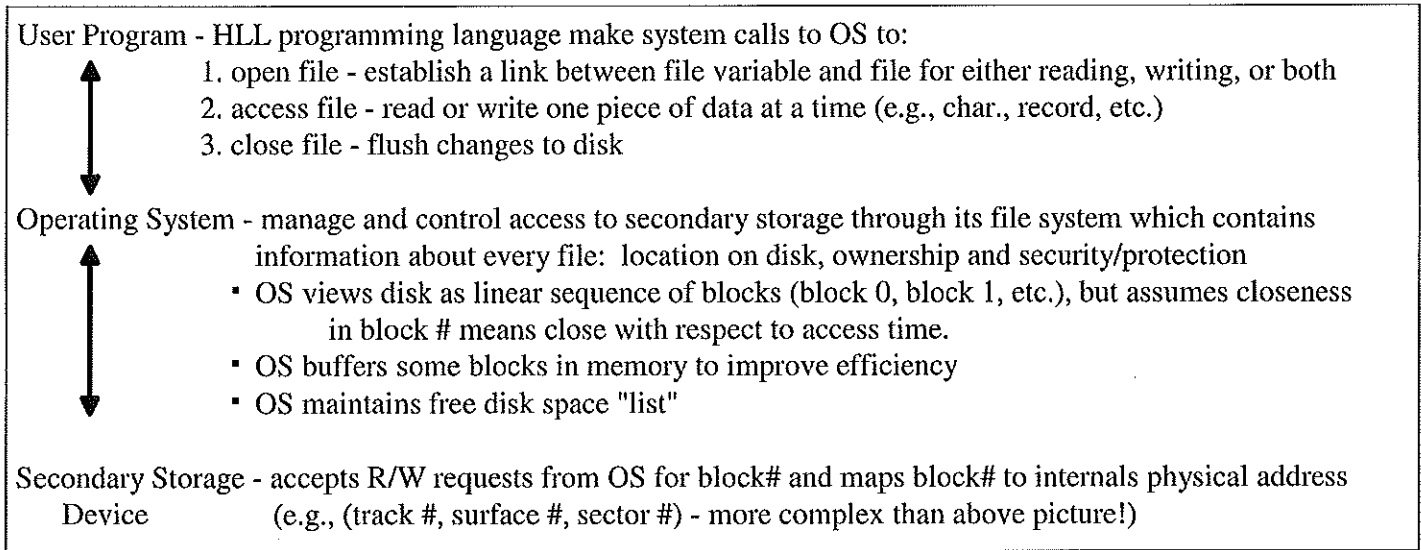
track #	surface #	sector #
---------	-----------	----------

3. Disk-access time = (seek time) + (rotational delay) + (data transfer time). How is each component of the disk-access time effected by increasing the-disk's RPMs (revolutions-per-minute)?

*rotational delay - on average about 1/2 rotation*  
*data transfer - time for sector to spin under R/W head*

b) If we want fast access to a collection of sectors, where can we place them to minimize seek time and rotational delay?

*consecutive sectors on same track (or cylinder)*




Kinds of File Access:

- serial/sequential files - open at the beginning and read sequentially from beginning to end linearly
- random-access files - "seek" to any position by specifying a byte-offset from the beginning of the file, record #, etc.
- random-access of a record by key

Implementation of Files on Disk- how are blocks allocated?


4. non-contiguous - scattered across linear address space of OS and disk



a) What types of file access are supported efficiently? *seq.*

b) How easy is it for the file to grow in size? *easy*

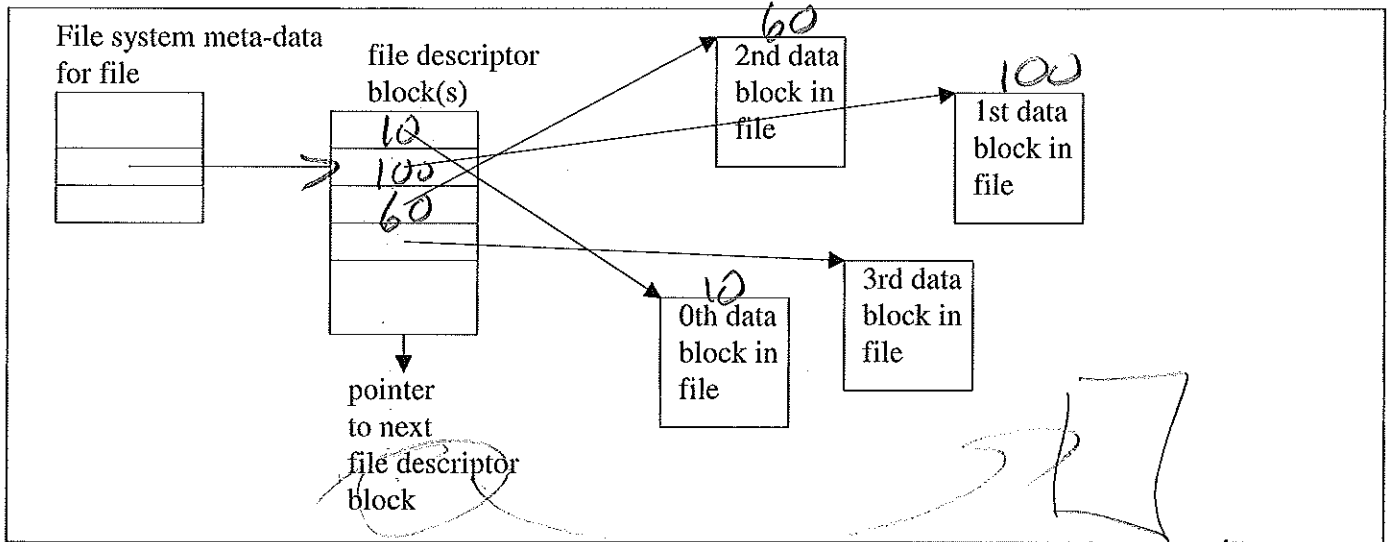
5. contiguous - sequential collection of blocks from OS linear view of disk



a) What types of file access are supported efficiently? *seq., random (seek), random (hashing) O(1)*

b) How easy is it for the file to grow in size? *hard*

6. file descriptor blocks - list of blocks hold the address of the physical location of data blocks




- a) What types of file access are supported efficiently? *seq., seek (random), random by key*
- b) How easy is it for the file to grow in size? *easy*

7. To implement "random-access of a record by key" in a file how might we use hashing?

*Using contiguous or file-descriptor block, we can run key through hash function to determine a "location" (home address) of key. Hopefully, we can read one disk sector and find target record.*

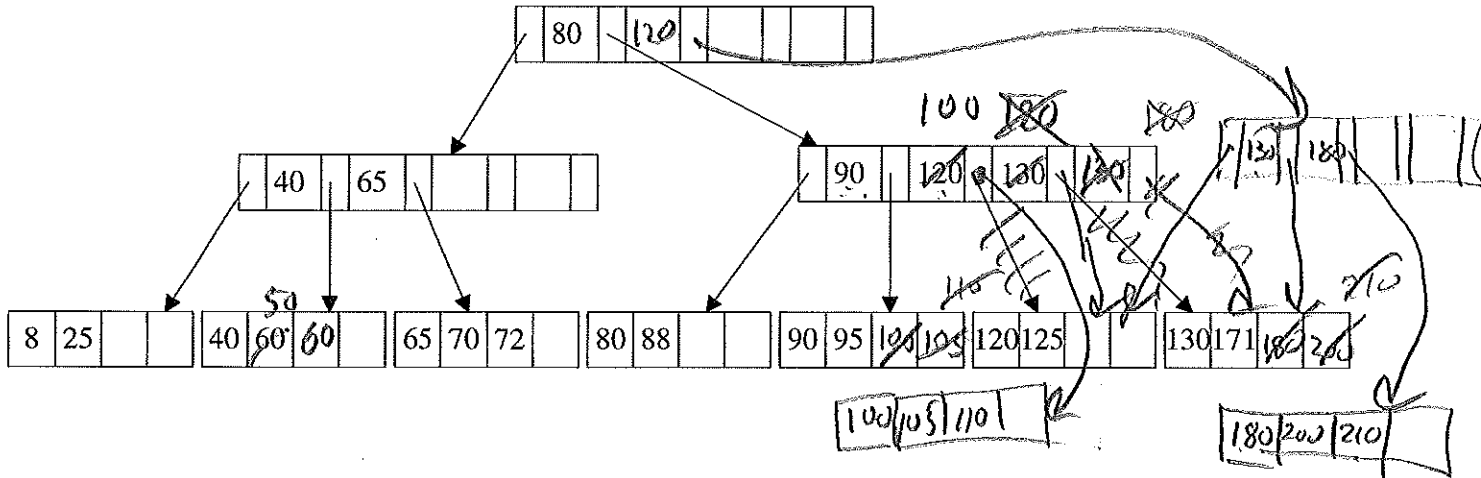
8. To implement "random-access of a record by key" in a file why would an AVL tree not work well?



9. A B+ Tree is a multi-way tree (typically in the order of 100s children per node) used primarily as a file-index structure to allow fast search (as well as insertions and deletions) for a target key on disk. Two types of pages (B+ tree "nodes") exist:

- Data pages - which always appear as leaves on the same level of a B+ tree (usually a doubly-linked list too)
- Index pages - the root and other interior nodes above the data page leaves. Index nodes contain some minimum and maximum number of keys and pointers bases on the B+ tree's *branching factor* ( $b$ ) and *fill factor*. A 50% fill factor would be the minimum for any B+ tree. All index pages must have  $\lceil b/2 \rceil \leq \# \text{ child} \leq b$ , except the root which must have at least two children.

Consider an B+ tree example with  $b = 5$ .



a) How would you find 88?

b) The insert algorithm for a B+ tree is summarized by the below table. Where would you insert 50, 100, 105, 110, 180, 200, 210?

Situation		insertion Algorithm
Data Page Full?	Parent Index Page Full?	
No	No	Place record in sorted position in the appropriate data page.
Yes	No	<ol style="list-style-type: none"> <li>1. Split data page with records <math>&lt;</math> middle key going in left data page and records <math>\geq</math> middle key going in right data page.</li> <li>2. Place middle key in index page in sorted order with the pointer immediately to its left pointing to the left data page and the pointer immediately to its right pointing to the right data page.</li> </ol>
Yes	Yes	<ol style="list-style-type: none"> <li>1. Split data page with records <math>&lt;</math> middle key going in left data page and records <math>\geq</math> middle key going in right data page.</li> <li>2. Adding middle key to parent index page causes it to split with keys <math>&lt;</math> middle key going into the left index page, keys <math>&gt;</math> middle key going in right index page, <b>and</b> the middle key inserted into the next higher level index page. If the next higher index page is full continue to splitting index pages up the B+ tree as necessary.</li> </ol>