

CARNEGIE MELLON UNIVERSITY  
COMPUTER SCIENCE DEPARTMENT  
15-445/645 – DATABASE SYSTEMS (FALL 2022)  
PROF. ANDY PAVLO

Homework #2 (by Mike Xu) – Solutions  
Due: **Sunday September 25, 2022 @ 11:59pm**

**IMPORTANT:**

- Enter all of your answers into **Gradescope by 11:59pm on Sunday September 25, 2022.**
- **Plagiarism:** Homework may be discussed with other students, but all homework is to be completed **individually**.

For your information:

- Graded out of **100** points; **4** questions total
- Rough time estimate:  $\approx$ 1-4 hours (0.5-1 hours for each question)

*Revision : 2022/09/30 14:36*

Question	Points	Score
Storage Models	16	
Cuckoo Hashing	20	
Extendible Hashing	28	
B+Tree	36	
Total:	100	

**Question 1: Storage Models.....[16 points]****Graded by:**

Consider a database with a single table  $R(q\_id, txns, total, failed)$ , where  $q\_id$  is the *primary key*, and all attributes are the same fixed width. Suppose  $R$  has 20,000 tuples that fit into 100 pages, Ignore any additional storage overhead for the table (e.g., page headers, tuple headers). Additionally, you should make the following assumptions:

- The DBMS does *not* have any additional meta-data (e.g., sort order, zone maps).
- $R$  does *not* have any indexes (including for primary key  $q\_id$ )
- None of  $R$ 's pages are already in the buffer pool.

Consider the following query:

```
SELECT total - failed FROM R
WHERE q_id = 96 AND txns > 420;
```

- (a) Suppose the DBMS uses the decomposition storage model (DSM) with implicit offsets
- i. [4 points] What is the *minimum* number of pages that the DBMS will potentially have to read from disk to answer this query?
- 1    2-10    11-50    51-100     $\geq 101$   
 Not possible to determine

**Solution:** 4 pages. 1 to find the primary key, + 3 to access  $txns$ ,  $total$ ,  $failed$  at their corresponding offsets.

- ii. [4 points] What is the *maximum* number of pages that the DBMS will potentially have to read from disk to answer this query?
- 1    2-10    11-50    51-100     $\geq 101$   
 Not possible to determine

**Solution:** 28 pages. There are 25 pages per attribute. In the worst case, we scan through all 25 pages to find the primary key, and then + 3 to access  $txns$ ,  $total$ ,  $failed$  at their corresponding offsets.

- (b) Suppose the DBMS uses the N-ary storage model (NSM)
- i. [4 points] What is the *minimum* number of pages that the DBMS will potentially have to read from disk to answer this query?
- 1    2-10    11-50    51-100     $\geq 101$   
 Not possible to determine

**Solution:** We find the tuple with the matching primary key on the first page. No need to look in other pages since all attributes are stored together.

- ii. [4 points] What is the *maximum* number of pages that the DBMS will potentially have to read from disk to answer this query?
- 1    2-10    11-50    51-100     $\geq 101$   
 Not possible to determine

**Solution:** We find the tuple with the matching primary key on the last page after scanning through every page.

**Question 2: Cuckoo Hashing..... [20 points]****Graded by:**

Consider the following cuckoo hashing schema:

1. Both tables have a size of 4.
2. The hashing function of the first table returns the fourth and third least significant bits:  
 $h_1(x) = (x \gg 2) \& 0b11$ .
3. The hashing function of the second table returns the least significant two bits:  
 $h_2(x) = x \& 0b11$ .
4. When inserting, try table 1 first.
5. When replacement is necessary, first select an element in the second table.
6. The original entries in the table are shown in the figure below.

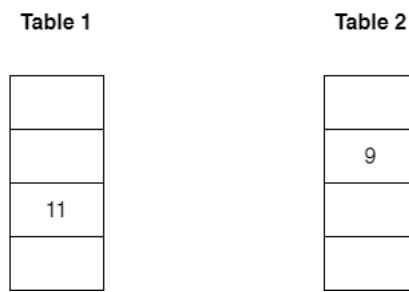


Figure 1: Initial contents of the hash tables.

- (a) [3 points] Select the sequence of insert operations that results in the initial state.
- Insert 9, insert 11     **Insert 11, insert 9**     None of the above

**Solution:** 11 is inserted into table 1  $0b10$  based on  $h_1$ , 9 experiences a collision and is hashed to table 2  $0b01$  based on  $h_2$ .

(b) [3 points] Insert key 16 and delete 11. Select the resulting two tables.

A)

Table 1	Table 2
	16
	9

C)

Table 1	Table 2
16	
	9

B)

Table 1	Table 2
16	
9	

D)

Table 1	Table 2
	16
9	

**Solution:** 16 is inserted into table 1  $0b00$  based on  $h_1$ .

(c) [4 points] Then insert 17 followed by 10. Select the resulting two tables.

A)

Table 1	Table 2
16	
	17
9	10

C)

Table 1	Table 2
17	16
	9
10	

B)

Table 1	Table 2
17	16
9	10

D)

Table 1	Table 2
17	16
	9
	10

**Solution:** 17 is inserted into table 2  $\emptyset b01$ , replacing 9, which is rehased into table 1  $\emptyset b10$ . 10 is inserted into table 2  $\emptyset b10$  because 9 is already in table 1  $\emptyset b10$ .

(d) [5 points] Finally, insert 33 and delete 16. Select the resulting two tables.

A)

Table 1	Table 2
33	
	9
10	
	17

C)

Table 1	Table 2
33	
	17
9	10

B)

Table 1	Table 2
17	
	33
9	10

D)

Table 1	Table 2
9	33
10	
	17

**Solution:** Inserting 33 into table 2  $0b01$  replaces 17. 17 is rehashed into table 1  $0b00$ , replacing 16. 16 is rehashed to table 2  $0b00$ , but is then deleted

(e) [5 points] What is the smallest key that potentially causes an infinite loop given the tables in (d)?

0    1    2    6    9    10    None of the above

**Solution:** 1, 17, and 33 all have  $h_1(x) = 0$ ,  $h_2(x) = 1$ , which would continue to replace each other.

**Question 3: Extendible Hashing.....[28 points]****Graded by:**

Consider an extendible hashing structure such that:

- Each bucket can hold up to two records.
- The hashing function uses the lowest  $g$  bits, where  $g$  is the global depth.

(a) Starting from an empty table, insert keys 15, 14, 23, 11, 9.

i. [4 points] What is the global depth of the resulting table?

- 0    1    2    3    4    None of the above

**Solution:** The global depth is the largest depth on buckets, which is  $d = 3$  on buckets  $b011$  and  $b111$

ii. [4 points] What is the local depth the bucket containing 15?

- 0    1    2    3    4    None of the above

**Solution:** 15 is hashed to bucket  $b111$  along with 23, which split after the insertion of 11

iii. [4 points] What is the local depth of the bucket containing 14?

- 0    1    2    3    4    None of the above

**Solution:** 14 is the only element hashed to bucket  $0$ , pointed to by  $b000$ ,  $b010$ ,  $b100$ , and  $b110$ .

(b) Starting from the result in (a), you insert keys 12, 5, 7, 13, 2.

i. [4 points] Which key will first cause a split (without doubling the size of the table)?

- 12    5    7    13    2    None of the above

**Solution:** Inserting 12 (into bucket  $b0$ ) and 5 (into bucket  $b01$ ) does not cause any splits. Inserting 7 (into bucket  $b111$ ) increases the global depth to 4, so it does not meet the criteria. Inserting 13 will split bucket  $b01$ .

ii. [4 points] Which key will first make the table double in size?

- 12    5    7    13    2    None of the above

**Solution:** See above.



(c) Now consider the table below, along with the following deletion rules:

1. If two buckets satisfy the following:
  - (a) They have the same local depth  $d$
  - (b) They share the first  $d - 1$  bits of their indexes (e.g.  $b010$  and  $b110$  share the first 2 bits)
  - (c) Their constituent elements fit in a single bucket.
 Then they can be merged into a single bucket with local depth  $d - 1$ .
2. If the global depth  $g$  becomes strictly greater than all local depths, then the table can be halved in size. The resulting global depth is  $g - 1$ .

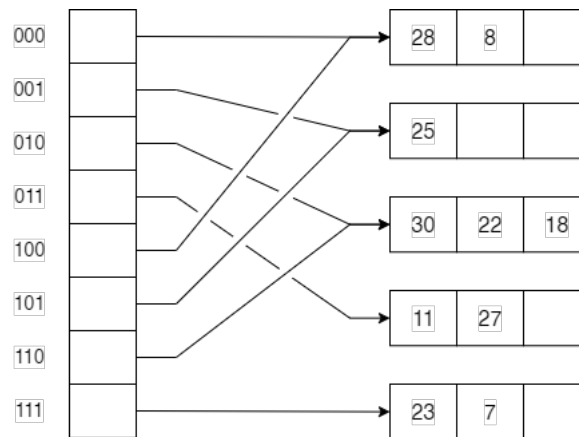


Figure 2: Extensible Hash Table along with the indexes of each bucket

Starting from the table above, delete keys 25, 18, 22, 27, 7.

- i. [4 points] Which deletion first causes a reduction in a local depth.
- 25    18    22    27    7    None of the above

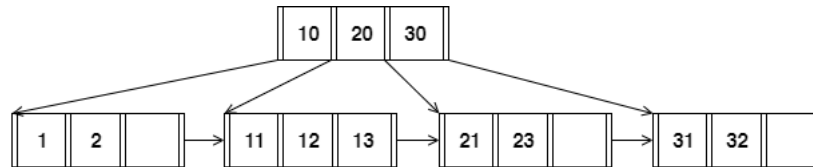
**Solution:** Deleting 25 and 18 does not make their bucket qualify for merging. Deleting 22 from bucket  $b10$  allows it to merge with  $b00$ .

- ii. [4 points] Which deletion first causes a reduction in global depth.
- 25    18    22    27    7    None of the above

**Solution:** Deleting 27 from bucket  $b011$  allows it to merge with  $b111$ . Since these two buckets are the only ones of depth  $d = 3$ , this merge reduces the global depth to  $d = 2$

**Question 4: B+ Tree** ..... [36 points]**Graded by:**

Consider the following B+ tree.

Figure 3: B+ Tree of order  $d = 4$  and height  $h = 2$ .

When answering the following questions, be sure to follow the procedures described in class and in your textbook. You can make the following assumptions:

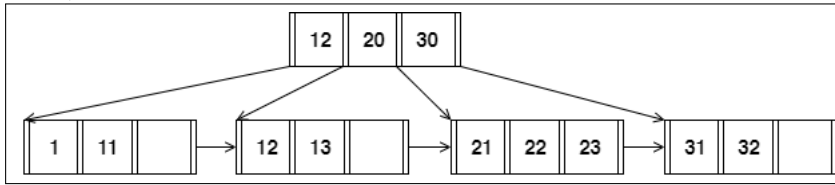
- A left pointer in an internal node guides towards keys  $<$  than its corresponding key, while a right pointer guides towards keys  $\geq$ .
  - A leaf node underflows when the number of **keys** goes below  $\lceil \frac{d-1}{2} \rceil$ .
  - An internal node underflows when the number of **pointers** goes below  $\lceil \frac{d}{2} \rceil$ .
- (a) [2 points] How many pointers (parent-to-child and sibling-to-sibling) do you chase to find all keys between  $9^*$  and  $19^*$ ?

2    3    4    5    6    7

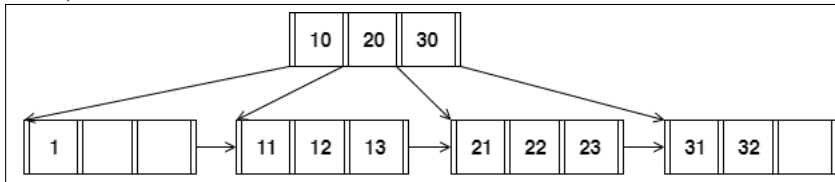
**Solution:** 1 to get to the leaf node. 2 more until it sees a key  $> 19^*$

(b) [6 points] Insert  $22^*$  into the B+tree, then delete  $2^*$ . Select the resulting tree.

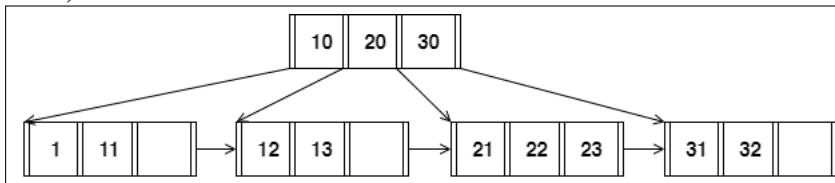
A)



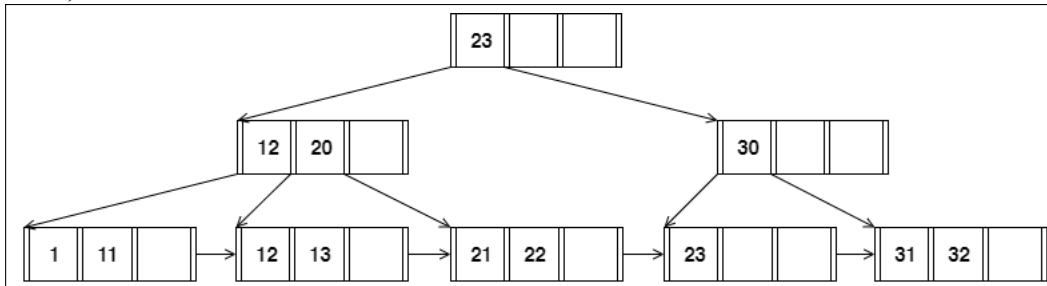
B)



C)



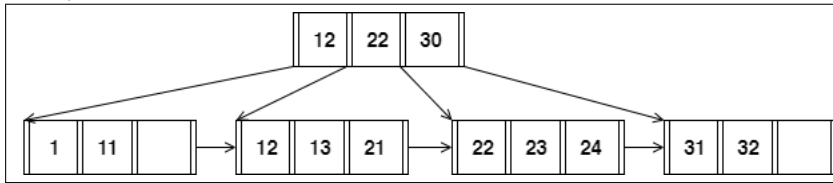
D)



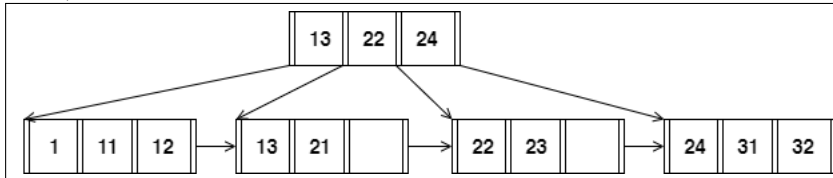
**Solution:** Insertion requires no splitting, deletion causes a redistribution.

(c) [10 points] Then Insert 24\*. Select the resulting tree.

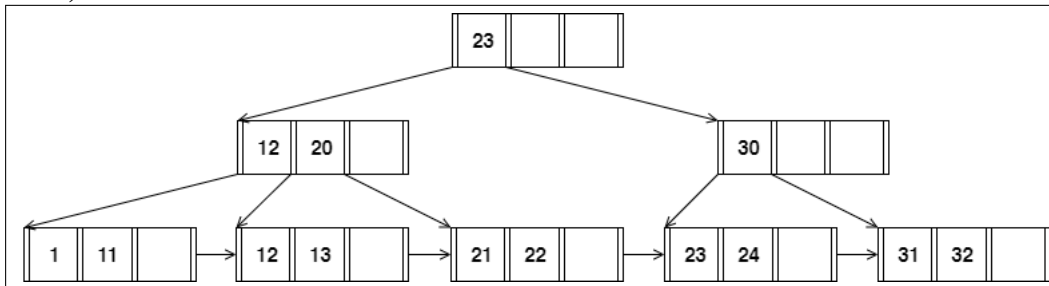
A)



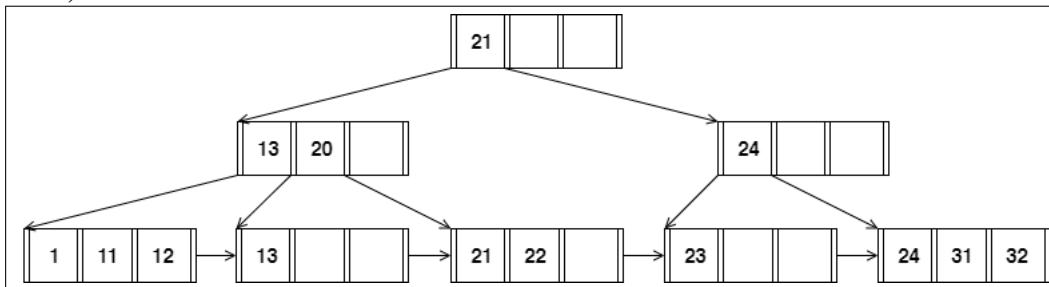
B)



C)



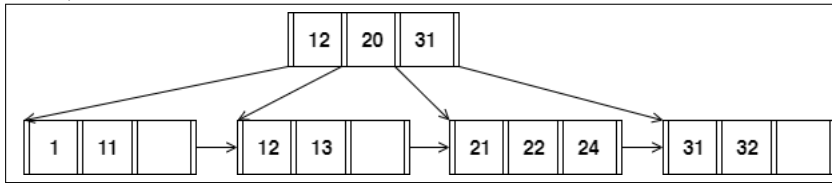
D)



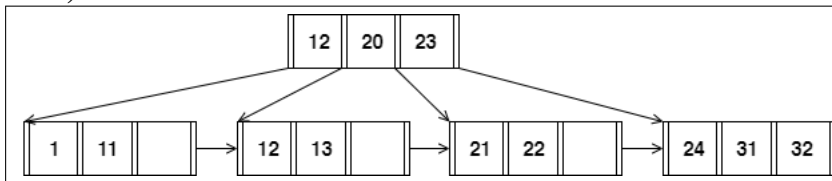
**Solution:** Insertion splits both the leaf node and the root/internal node.

(d) [10 points] Finally, delete 23\*. Select the resulting tree.

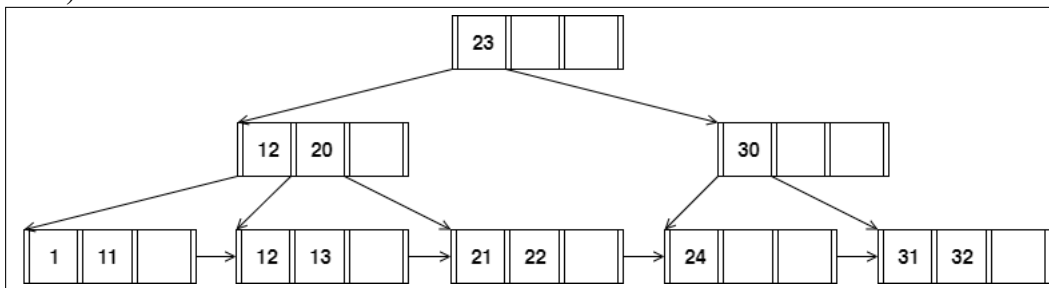
A)



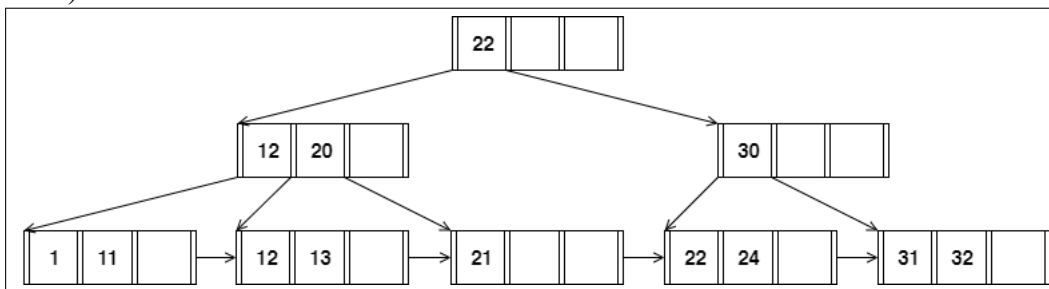
B)



C)



D)



**Solution:** Deletion merges the leaf nodes as well as the internal nodes. Root node is deleted and the merged internal node becomes the root.

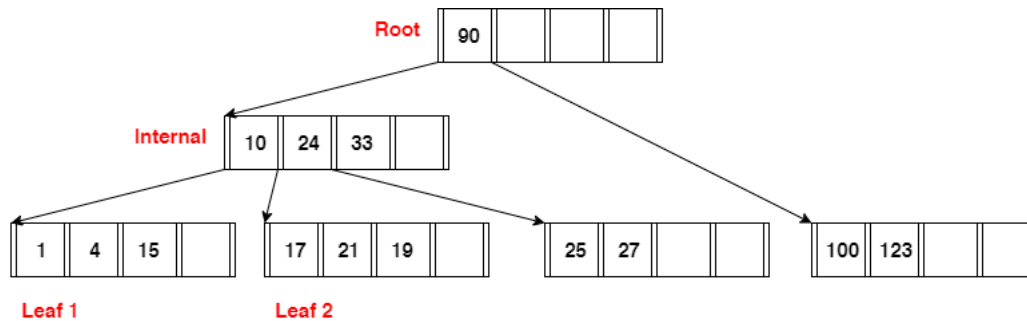


Figure 4: B+tree with violations

The B+Tree shown in Figure 4 is invalid. That is, its nodes violate the correctness properties of B+Trees that we discussed in class. If the tree is invalid, select all the properties that are violated for each node. If the node is valid, then select 'None'. There will be **no** partial credit for missing violations.

*Note: If a node's subtrees are not the same height, the balance property is violated at that node only.*

- i. [2 points] Which properties are violated by **Leaf 1**?
  - Key order property
  - Half-full property
  - Balance property
  - Separator keys
  - None**
- ii. [2 points] Which properties are violated by **Leaf 2**?
  - Key order property**
  - Half-full property
  - Balance property
  - Separator keys
  - None
- iii. [2 points] Which properties are violated by **Internal Node**?
  - Key order property
  - Half-full property
  - Balance property
  - Separator keys**
  - None
- iv. [2 points] Which properties are violated by **Root**?
  - Key order property
  - Half-full property
  - Balance property**
  - Separator keys
  - None