

CARNEGIE MELLON UNIVERSITY  
COMPUTER SCIENCE DEPARTMENT  
15-445/645 – DATABASE SYSTEMS (SPRING 2023)  
PROF. CHARLIE GARROD

Homework #2 (by Arvin Wu)  
Due: **Friday February 17, 2023 @ 11:59pm**

**IMPORTANT:**

- Enter all of your answers into **Gradescope by 11:59pm on Friday February 17, 2023.**
- **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:  $\approx$ 4-6 hours (1-1.5 hours for each question)

*Revision : 2023/02/07 15:29*

Question	Points	Score
Storage Models	16	
Cuckoo Hashing	27	
Extendible Hashing	22	
B+Tree	35	
Total:	100	

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

Consider a database with a single table  $T(\text{course\_id}, \text{course\_name}, \text{instructor}, \text{class\_size}, \text{hrs\_per\_week})$ , where  $\text{course\_id}$  is the *primary key*, and all attributes are the same fixed width. Suppose  $T$  has 5,000 tuples that fit into 500 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).
- $T$  does *not* have any indexes (including for primary key  $\text{course\_id}$ )
- None of  $T$ 's pages are already in the buffer pool.
- Content-wise, the tuples of  $T$  will make each query run the longest possible (this assumption is critical for solving part (a))
- The tuples of  $T$  can be in any order (this assumption is critical for solving part (b) when you compute the *minimum* versus *maximum* number of pages that the DBMS will potentially have to read)

(a) Consider the following query:

```
SELECT MAX(hrs_per_week) FROM T
WHERE class_size > 10;
```

- i. [3 points] Suppose the DBMS uses the decomposition storage model (DSM) with implicit offsets. How many pages will the DBMS potentially have to read from disk to answer this query? (Keep in mind our assumption about the contents of  $T$ !)
  - 1-100    101-200    201-300    301-500     $\geq 501$     Not possible to determine
- ii. [3 points] Suppose the DBMS uses the N-ary storage model (NSM). How many pages will the DBMS potentially have to read from disk to answer this query? (Keep in mind our assumption about the contents of  $T$ !)
  - 1-100    101-200    201-300    301-500     $\geq 501$     Not possible to determine

(b) Now consider the following query:

```
SELECT course_name, instructor, class_size FROM T
WHERE course_id = 15445 OR course_id = 15645;
```

- i. Suppose the DBMS uses the decomposition storage model (DSM) with implicit offsets.
  - $\alpha$ ) [3 points] What is the *minimum* number of pages that the DBMS will potentially have to read from disk to answer this query?
    - 1    2-4    5-100    101-200    201-500     $\geq 501$

- Not possible to determine
- $\beta$ ) **[3 points]** What is the *maximum* number of pages that the DBMS will potentially have to read from disk to answer this query?
- 1     2-4     5-100     101-200     201-500      $\geq 501$
- Not possible to determine
- ii. Suppose the DBMS uses the N-ary storage model (NSM).
- $\alpha$ ) **[2 points]** What is the *minimum* number of pages that the DBMS will potentially have to read from disk to answer this query?
- 1     2-4     5-100     101-200     201-500      $\geq 501$
- Not possible to determine
- $\beta$ ) **[2 points]** What is the *maximum* number of pages that the DBMS will potentially have to read from disk to answer this query?
- 1     2-4     5-100     101-200     201-500      $\geq 501$
- Not possible to determine

**Question 2: Cuckoo Hashing.....[27 points]**

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.

Table 1	Table 2
16	
	3

Figure 1: Initial contents of the hash tables.

- (a) **[1 point]** Select the sequence of insert operations that results in the initial state.  
 Insert 16, insert 3     Insert 3, insert 16     None of the above

(b) Insert key 8 and then delete 16. Select the value in each entry of the resulting two tables.

i. Table 1

$\alpha$ ) [1 point] Entry 0 (0b00)  3  8  Empty

$\beta$ ) [1 point] Entry 1 (0b01)  3  8  Empty

$\gamma$ ) [1 point] Entry 2 (0b10)  3  8  Empty

$\delta$ ) [1 point] Entry 3 (0b11)  3  8  Empty

ii. Table 2

$\alpha$ ) [1 point] Entry 0 (0b00)  3  8  Empty

$\beta$ ) [1 point] Entry 1 (0b01)  3  8  Empty

$\gamma$ ) [1 point] Entry 2 (0b10)  3  8  Empty

$\delta$ ) [1 point] Entry 3 (0b11)  3  8  Empty

(c) After the changes from part (b), insert key 27 and then insert 4. Select the value in each entry of the resulting two tables.

i. Table 1

$\alpha$ ) [1 point] Entry 0 (0b00)  3  8  27  4  Empty

$\beta$ ) [1 point] Entry 1 (0b01)  3  8  27  4  Empty

$\gamma$ ) [1 point] Entry 2 (0b10)  3  8  27  4  Empty

$\delta$ ) [1 point] Entry 3 (0b11)  3  8  27  4  Empty

ii. Table 2

$\alpha$ ) [1 point] Entry 0 (0b00)  3  8  27  4  Empty

$\beta$ ) [1 point] Entry 1 (0b01)  3  8  27  4  Empty

$\gamma$ ) [1 point] Entry 2 (0b10)  3  8  27  4  Empty

$\delta$ ) [1 point] Entry 3 (0b11)  3  8  27  4  Empty

(d) After the changes from parts (b) and (c), insert key 19 and then delete 27. Select the value in each entry of the resulting two tables.

i. Table 1

$\alpha$ ) [1 point] Entry 0 (0b00)  3  8  4  19  Empty

$\beta$ ) [1 point] Entry 1 (0b01)  3  8  4  19  Empty

$\gamma$ ) [1 point] Entry 2 (0b10)  3  8  4  19  Empty

$\delta$ ) [1 point] Entry 3 (0b11)  3  8  4  19  Empty

ii. Table 2

$\alpha$ ) [1 point] Entry 0 (0b00)  3  8  4  19  Empty

$\beta$ ) [1 point] Entry 1 (0b01)  3  8  4  19  Empty

$\gamma$ ) [1 point] Entry 2 (0b10)  3  8  4  19  Empty

$\delta$ ) [1 point] Entry 3 (0b11)  3  8  4  19  Empty

(e) [2 points] What is the smallest key that potentially causes an infinite loop given the table that results from part (d)?

33  34  35  36  37  51  None of the above

**Question 3: Extendible Hashing.....[22 points]**

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 new extendible hashing structure is initialized with  $g = 0$  and one empty bucket
- (a) Starting from an empty table, insert keys 6, 1.
- i. [1 point] What is the global depth of the resulting table?  
 0    1    2    3    4    None of the above
  - ii. [1 point] What is the local depth of the bucket containing 6?  
 0    1    2    3    4    None of the above
- (b) Starting from the result in (a), you insert keys 17, 10.
- i. [2 points] What is the global depth of the resulting table?  
 0    1    2    3    4    None of the above
  - ii. [1 point] What is the local depth of the bucket containing 10?  
 0    1    2    3    4    None of the above
- (c) Starting from the result in (b), you insert key 25.
- i. [3 points] What is the global depth of the resulting table?  
 0    1    2    3    4    None of the above
  - ii. [3 points] What is the local depth of the bucket containing 25?  
 0    1    2    3    4    None of the above
  - iii. [2 points] What is the local depth of the bucket containing 17?  
 0    1    2    3    4    None of the above
  - iv. [1 point] What is the local depth of the bucket containing 6?  
 0    1    2    3    4    None of the above
  - v. [1 point] Which value, if inserted, will hash to the same bucket as the bucket containing key 10?  
 4    7    9    13    None of the above
- (d) [2 points] Starting from the result in (c), which **key(s)**, if inserted next, will cause a split that doubles the table's size?  
 5    12    24    29    33    None of the above
- (e) [2 points] Starting from the result in (c), which **key(s)**, if inserted next, will cause a split **without** doubling the table's size?  
 5    12    24    29    33    None of the above
- (f) [3 points] Starting from an empty table, insert keys 32, 64, 128, 256. What is the global depth of the resulting table?  
 4    5    6    7    8     $\geq 9$

**Question 4: B+Tree.....[35 points]**

Consider the following B+tree.

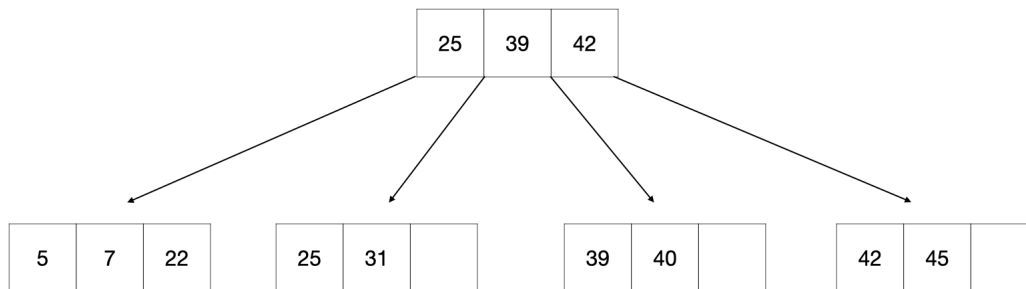


Figure 2: B+ Tree of order  $d = 4$  and height  $h = 2$ .

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

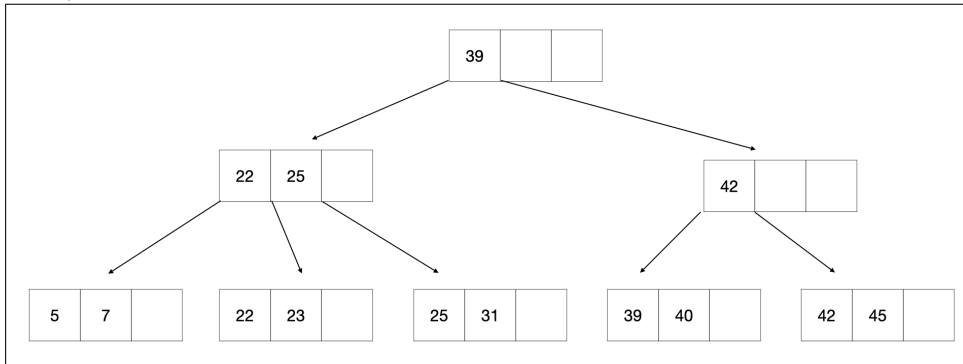
- A left pointer in an internal node guides towards keys  $<$  than its corresponding key, while a right pointer guides towards keys  $\geq$ .
- A leaf node underflows when the number of **keys** goes below  $\lceil \frac{d-1}{2} \rceil$ .
- An internal node underflows when the number of **pointers** goes below  $\lceil \frac{d}{2} \rceil$ .

Note that B+ tree diagrams for this problem omit leaf pointers for convenience. The leaves of actual B+ trees are linked together via pointers, forming a singly linked list allowing for quick traversal through all keys.

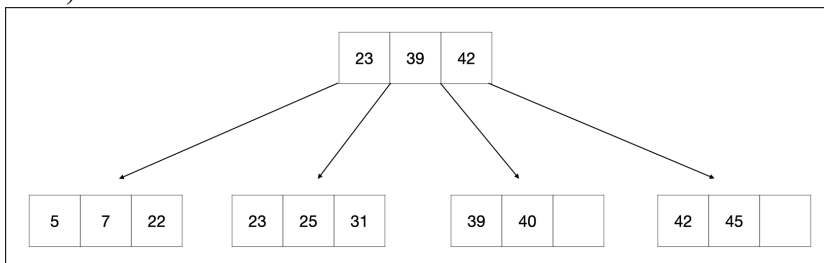


(a) [5 points] Insert  $23^*$  into the B+tree. Select the resulting tree.

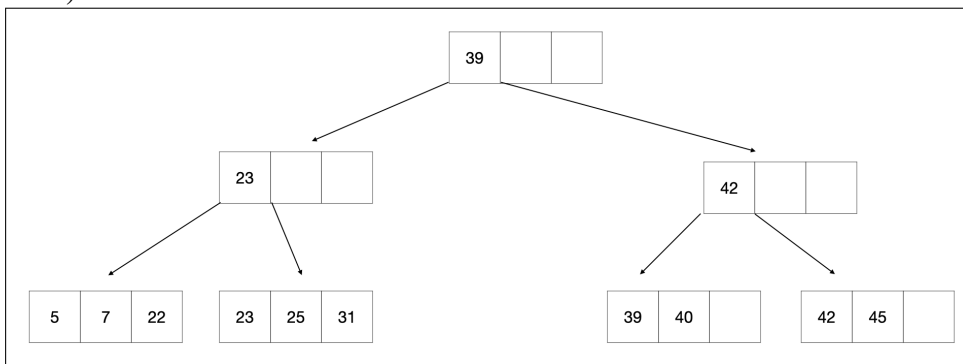
A)



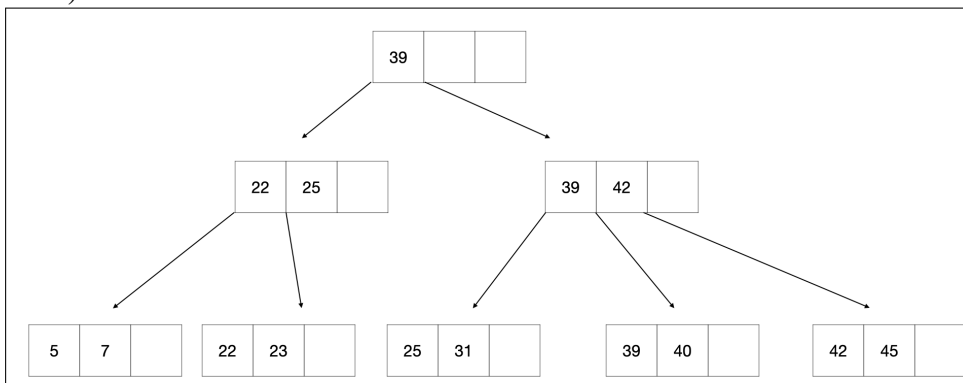
B)



C)

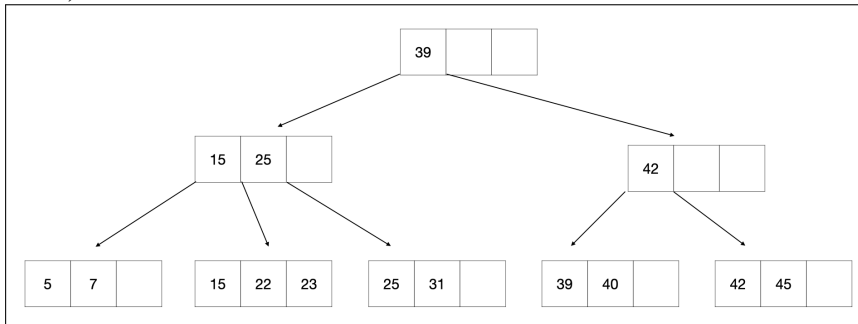


D)

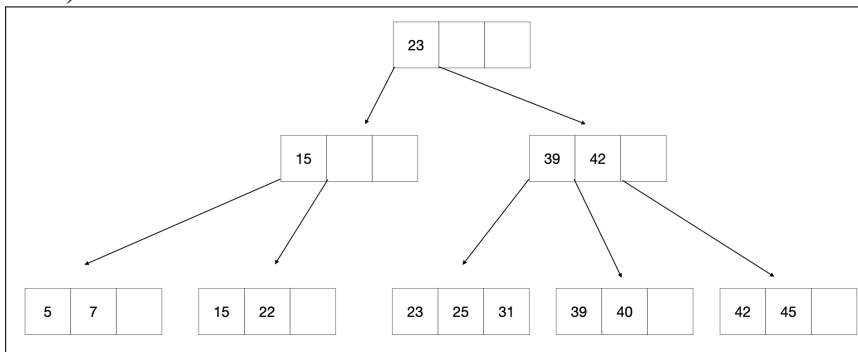


(b) [5 points] Starting with the tree that results from (a), insert 15\*. Select the resulting tree.

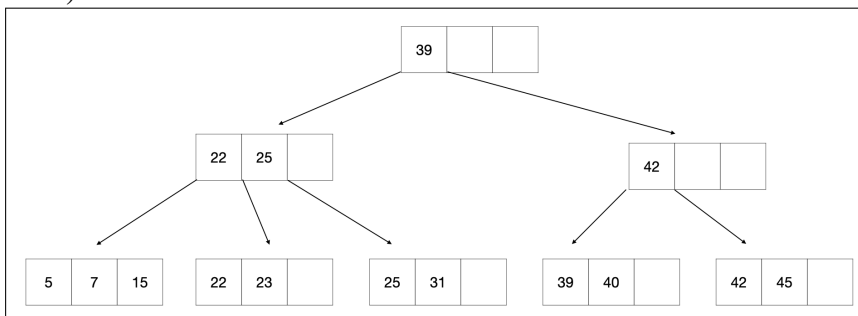
A)



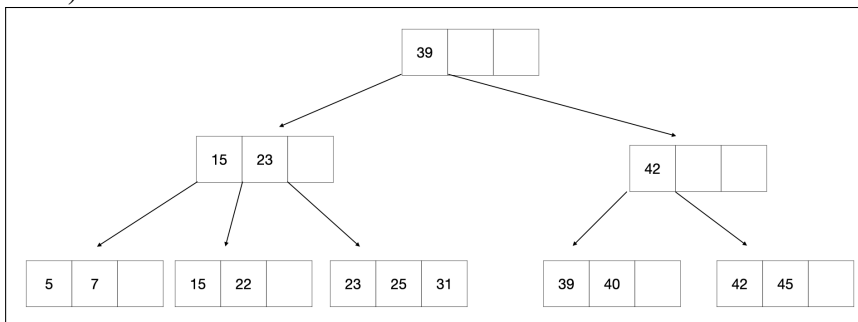
B)



C)

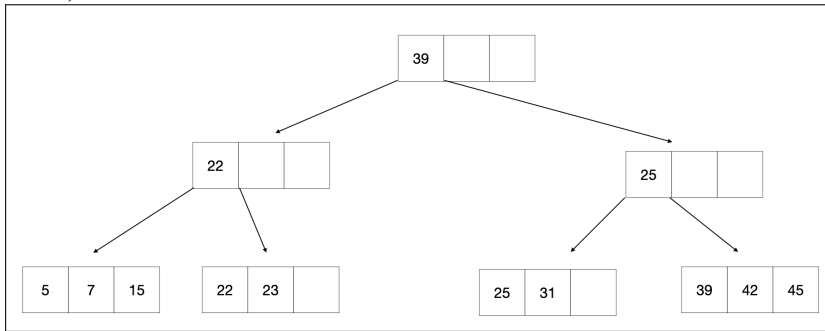


D)

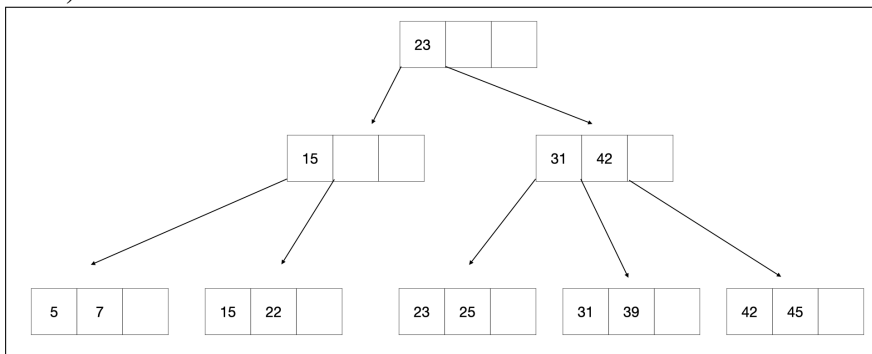


(c) [5 points] Starting with the tree that results from (b), delete 40\*. Select the resulting tree.

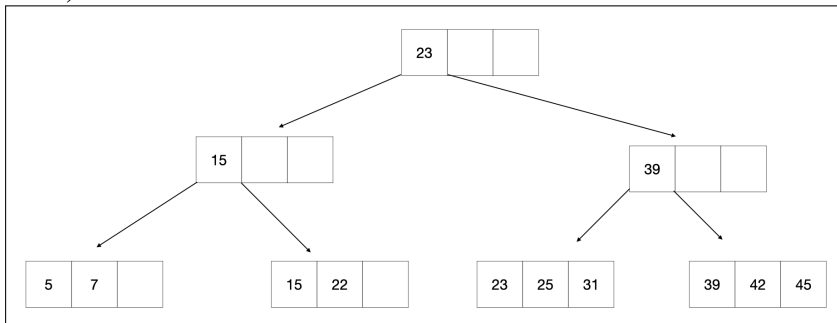
A)



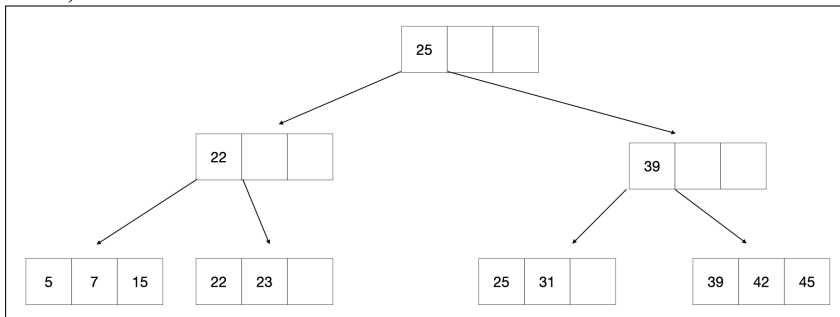
B)



C)

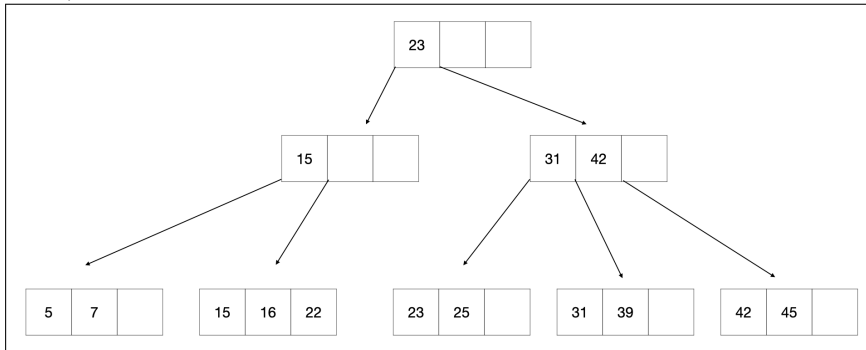


D)

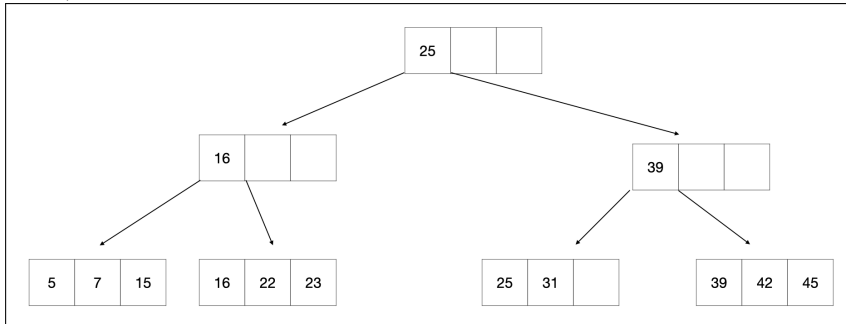


(d) [5 points] Starting with the tree that results from (c), insert 16\*. Select the resulting tree.

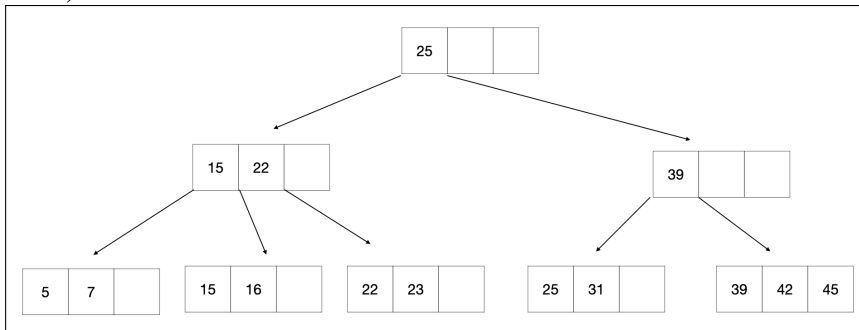
A)



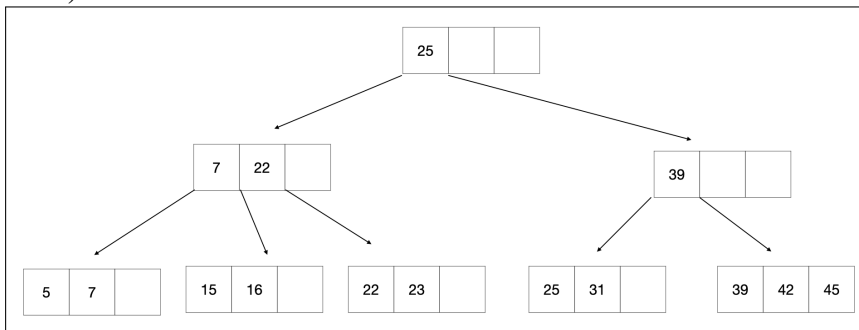
B)



C)

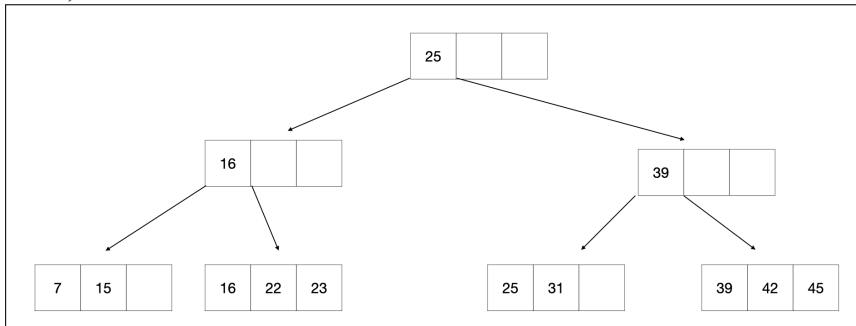


D)

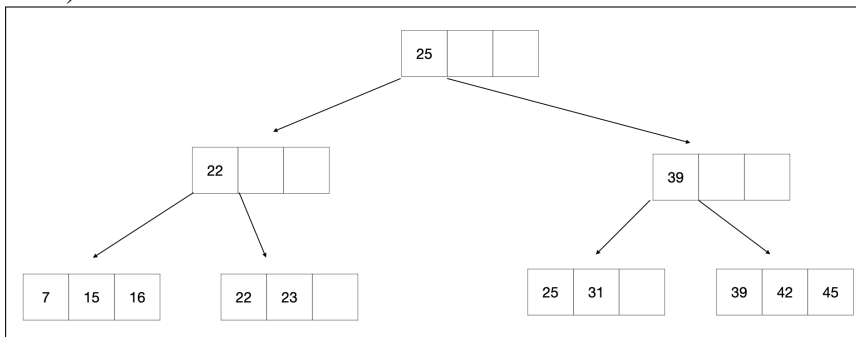


(e) [5 points] Finally, starting with the tree that results from (d), delete 5\*. Select the resulting tree.

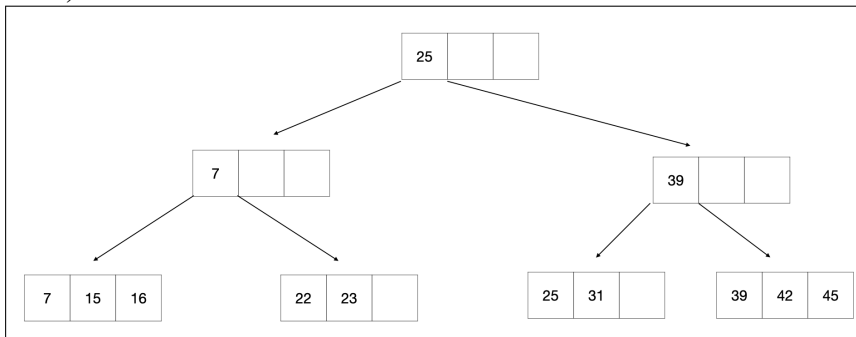
A)



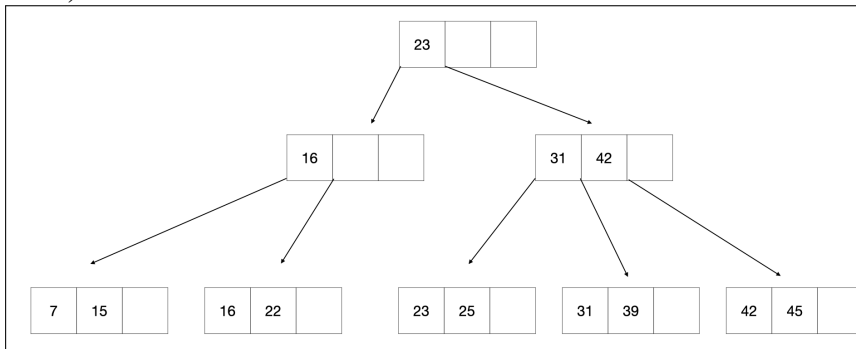
B)



C)



D)



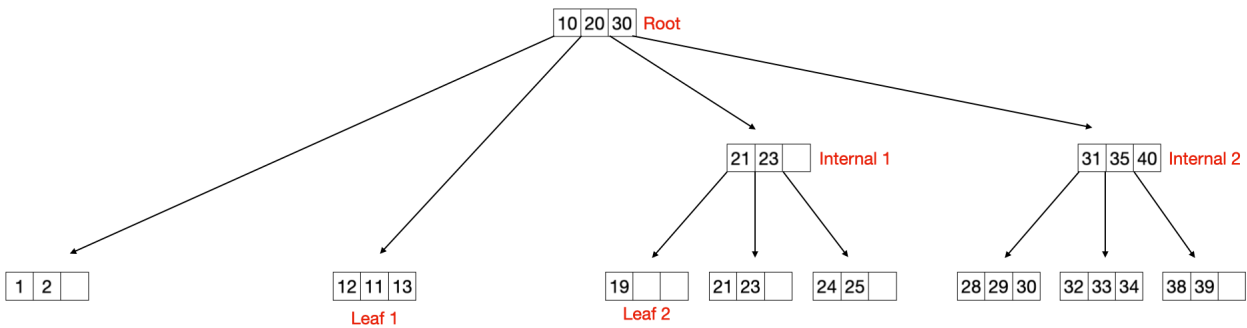


Figure 3: B+tree with violations

- (f) The B+Tree shown in Figure 3 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.
- If a node's subtrees contain values not in the range specified by the node's separator keys, the separator keys property is violated at that node.

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