





# Today's Agenda

---

Recap

Latches Overview

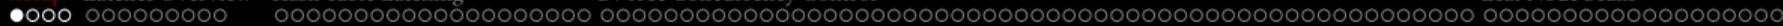
Hash Table Latching

B+Tree Concurrency Control

Leaf Node Scans

$B^{link}$ -Tree

Conclusion

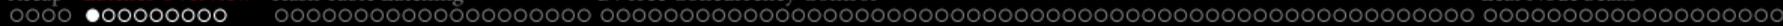


# Recap









# Latches Overview











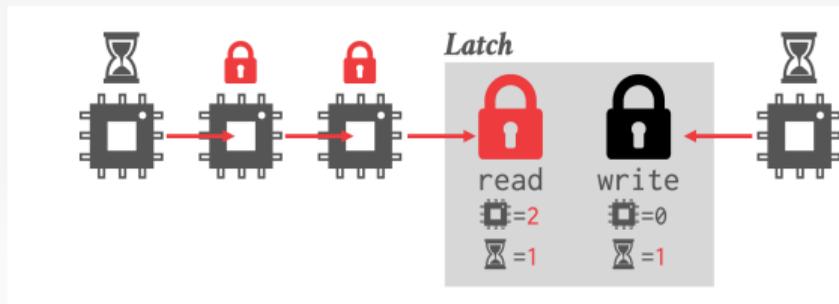




# Latch Implementations

- Approach 3: Reader-Writer Latch

- ▶ Allows for concurrent readers
- ▶ Must manage read/write queues to avoid starvation
- ▶ Can be implemented on top of spinlocks







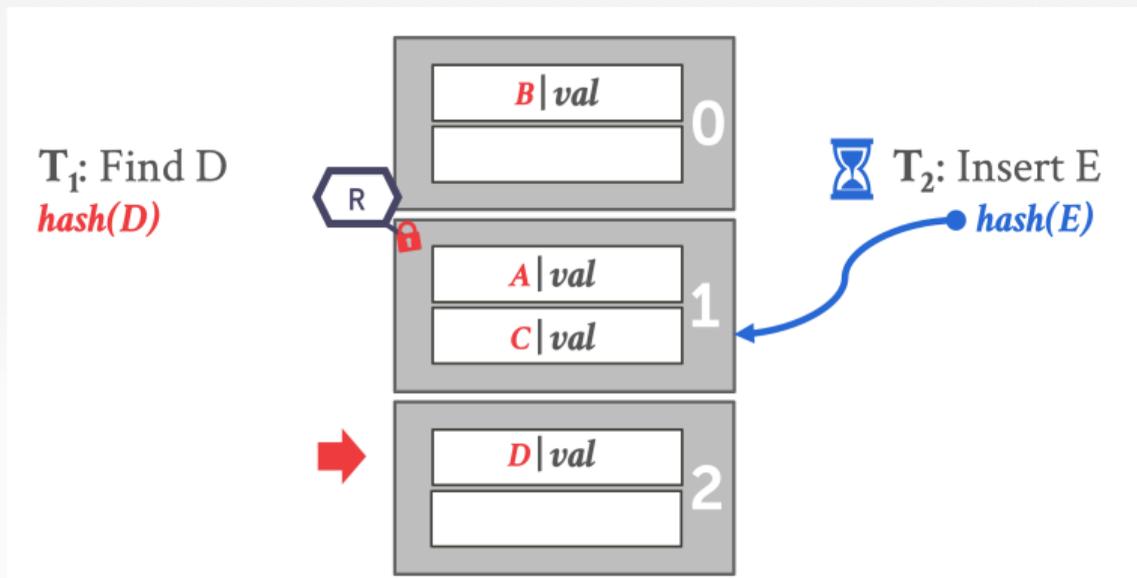








# Hash Table - Page Latches







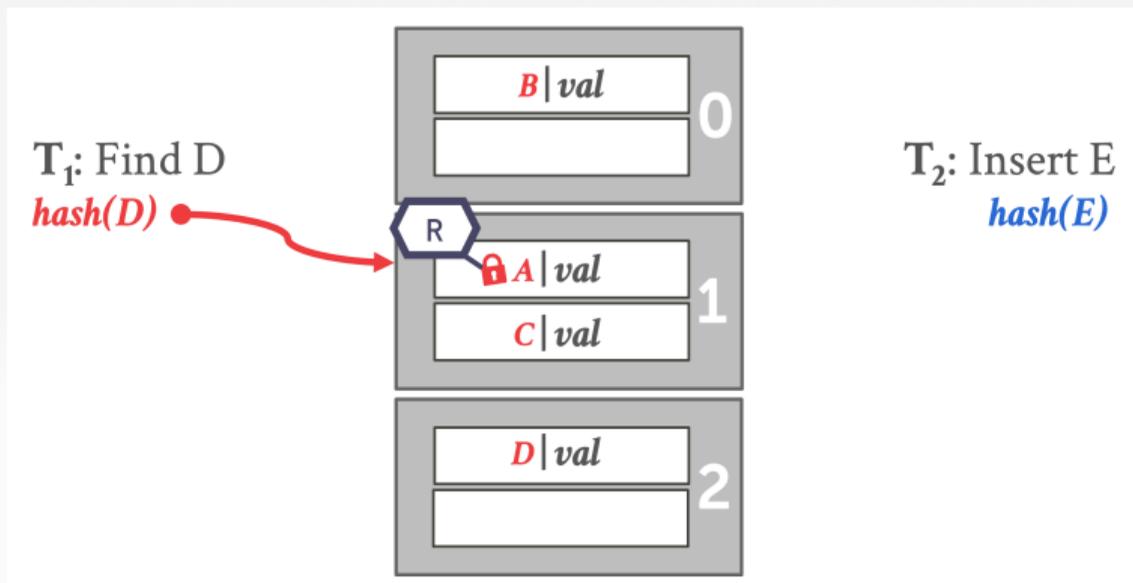






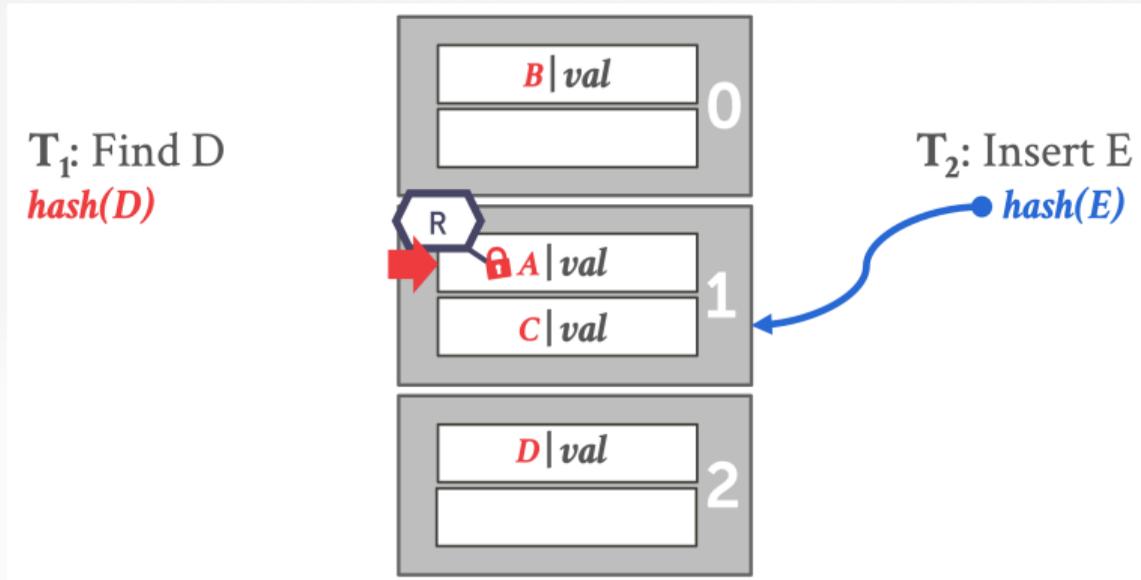


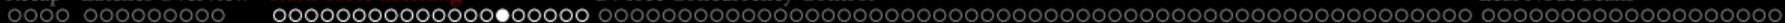
# Hash Table - Slot Latches



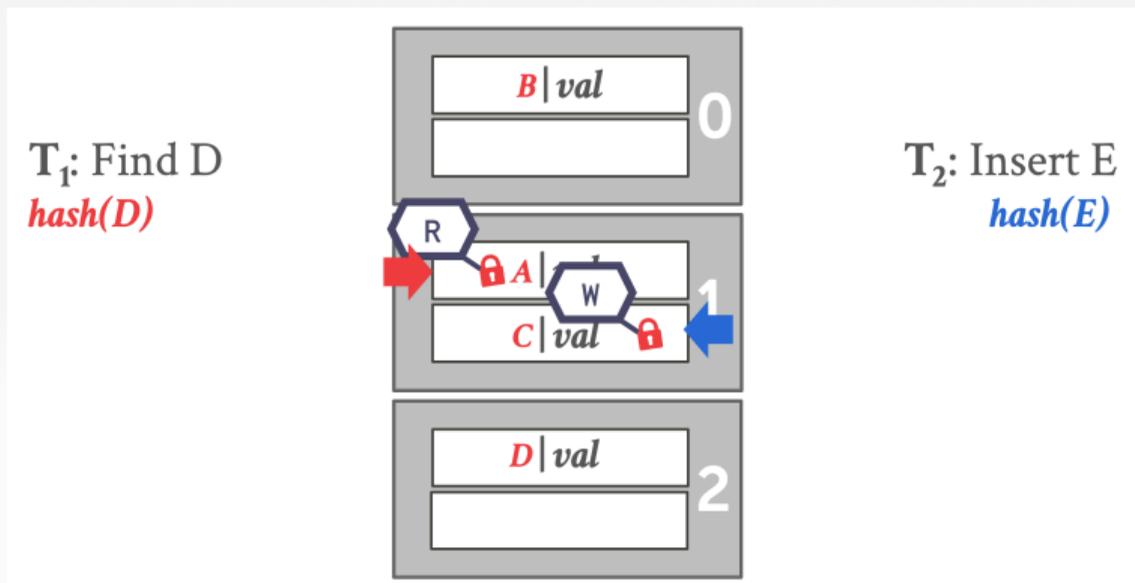


# Hash Table - Slot Latches



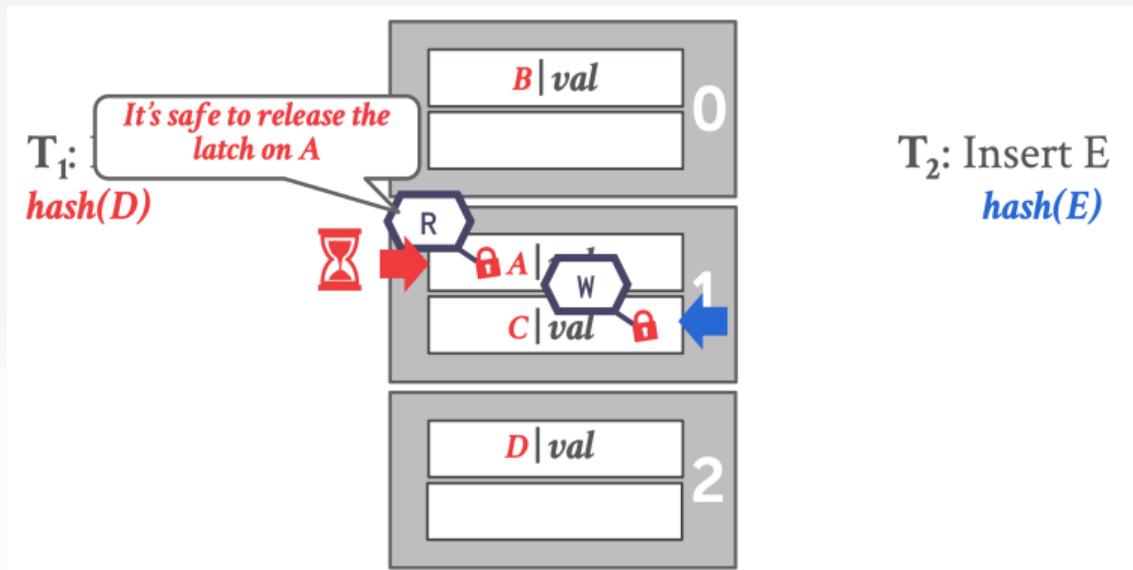


# Hash Table - Slot Latches





# Hash Table - Slot Latches

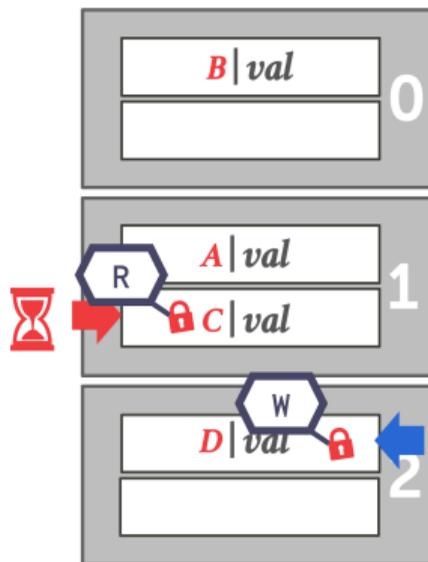






# Hash Table - Slot Latches

$T_1$ : Find D  
*hash(D)*

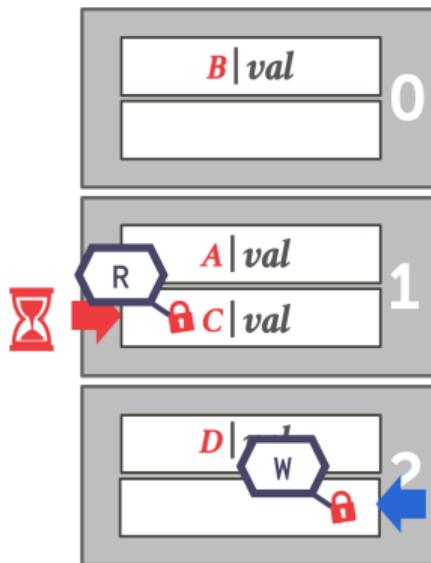


$T_2$ : Insert E  
*hash(E)*



# Hash Table - Slot Latches

$T_1$ : Find D  
*hash(D)*



$T_2$ : Insert E  
*hash(E)*





# B+Tree Concurrency Control

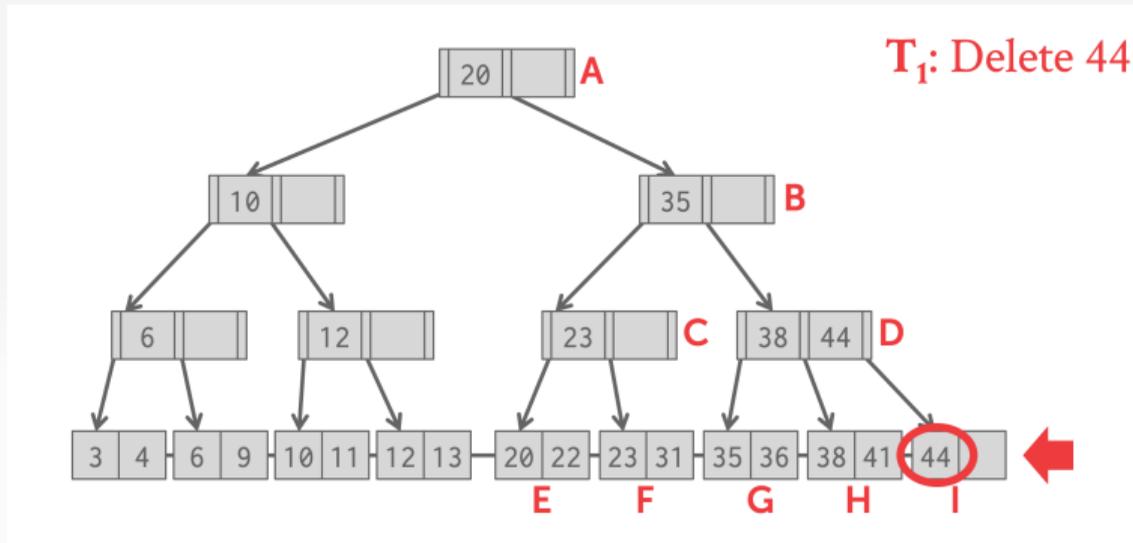
---

- We want to allow multiple threads to read and update a B+Tree at the same time.
- We need to handle two types of problems:
  - ▶ Threads trying to modify the contents of **a node** at the same time.
  - ▶ One thread traversing the tree while another thread splits/merges nodes.

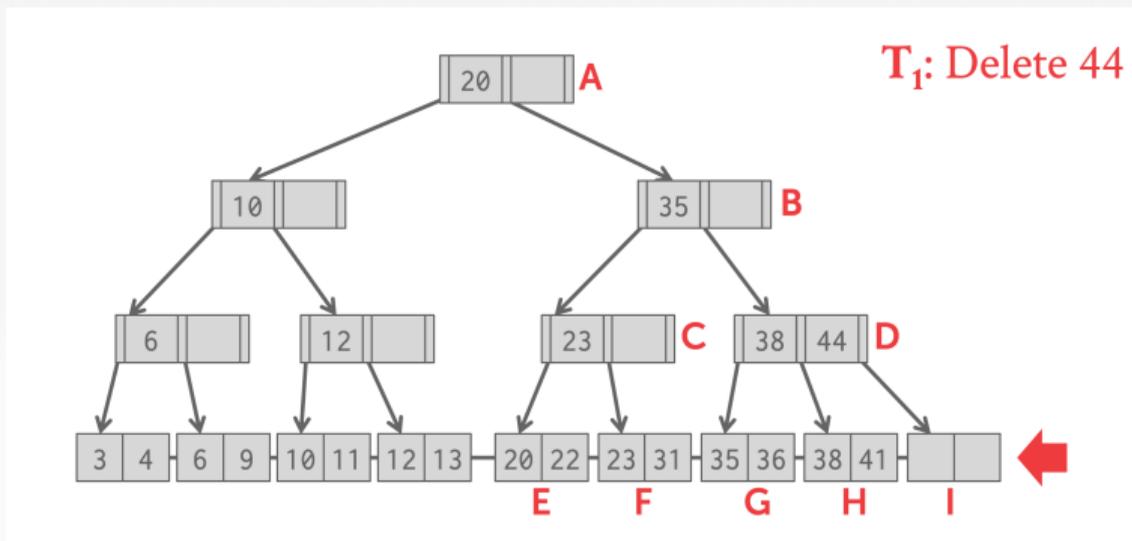




# B+Tree Concurrency Control: Example

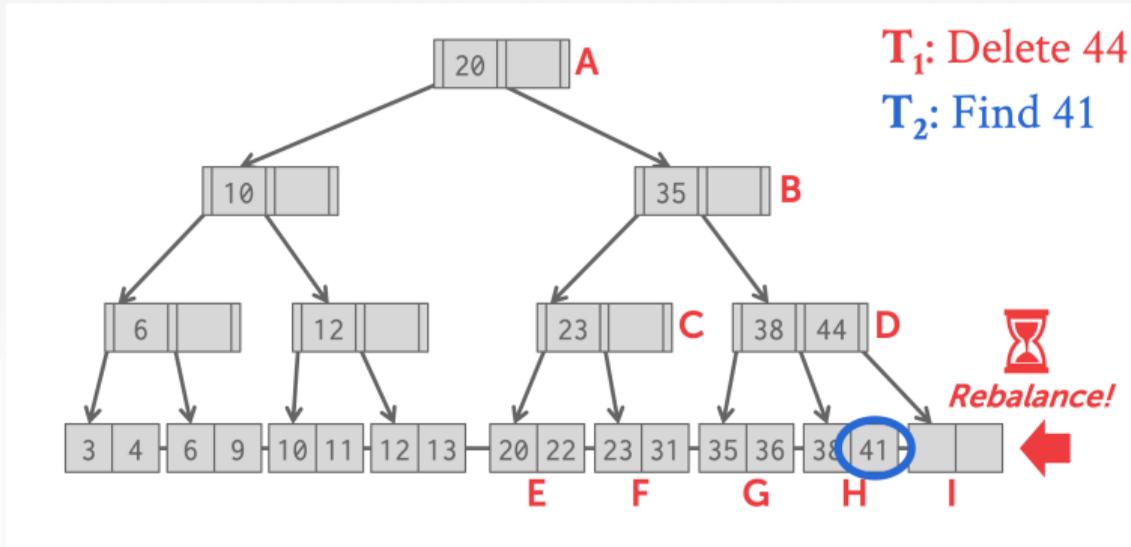


# B+Tree Concurrency Control: Example



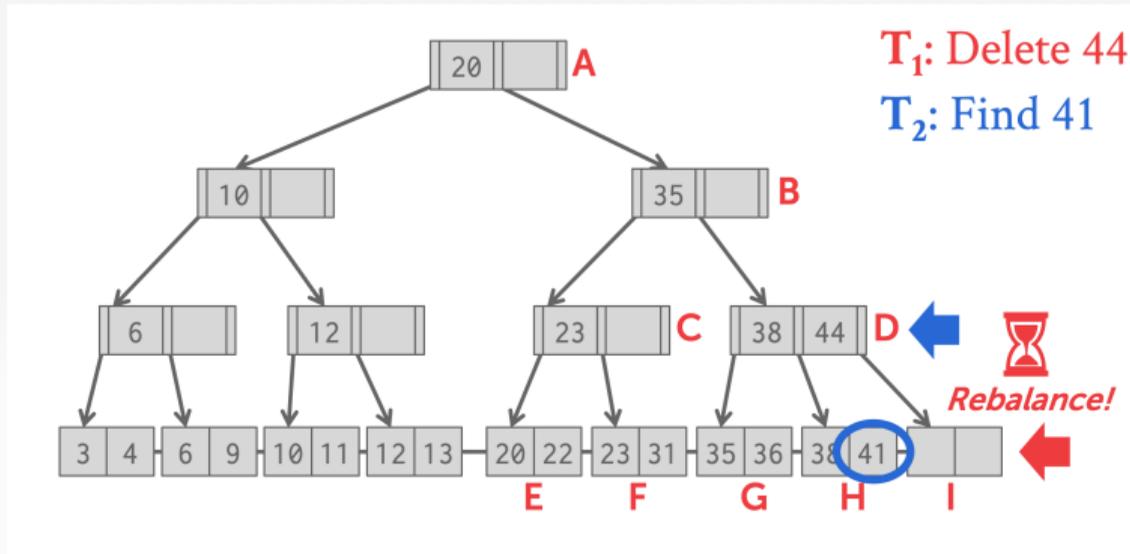


# B+Tree Concurrency Control: Example

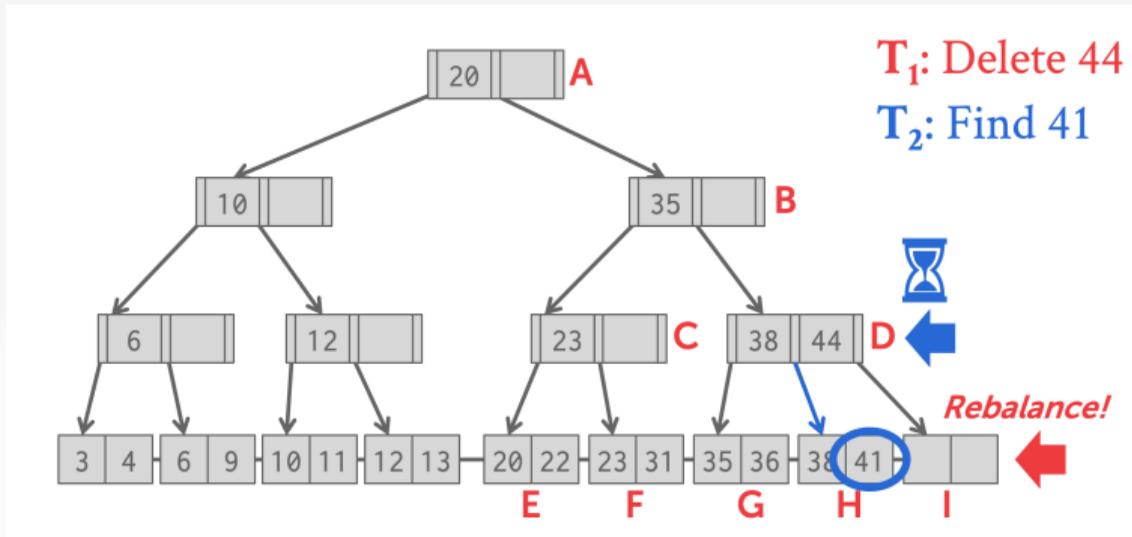




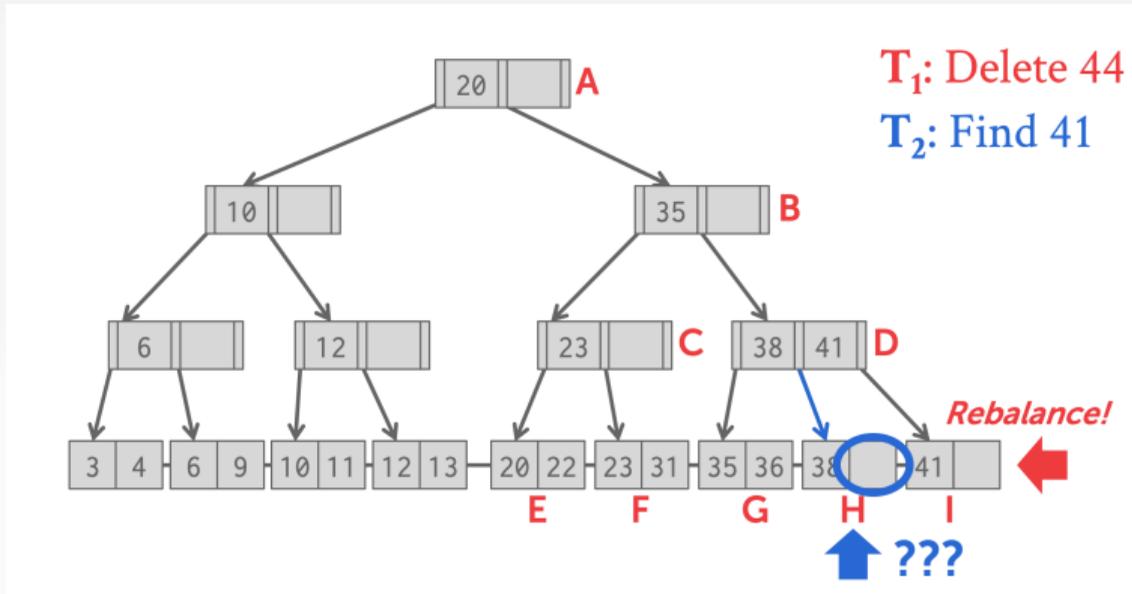
# B+Tree Concurrency Control: Example



# B+Tree Concurrency Control: Example



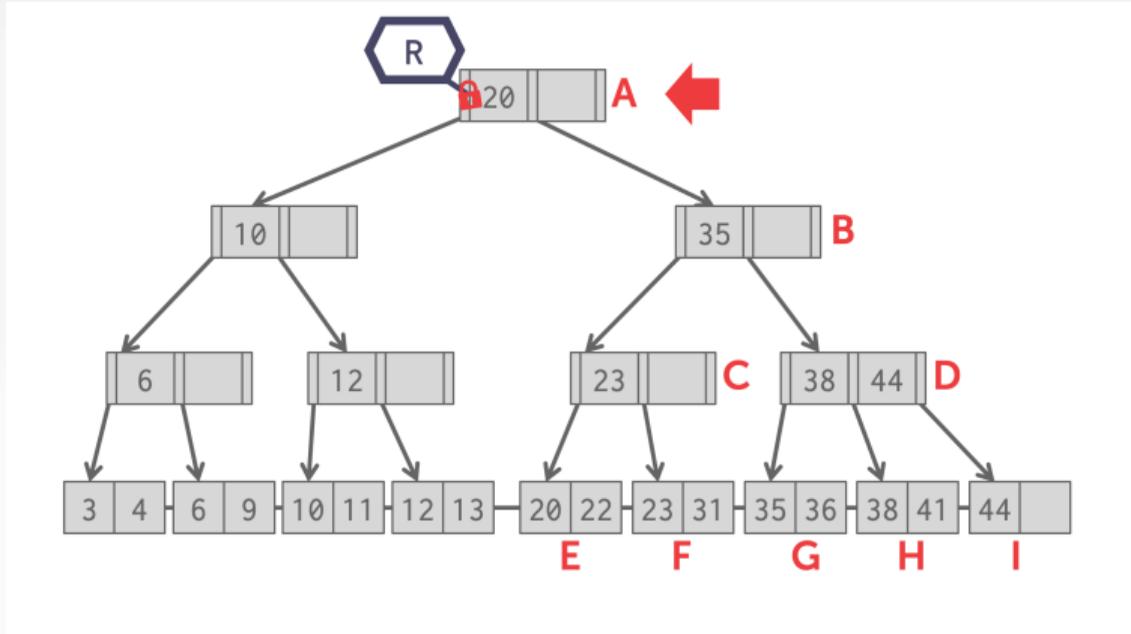
# B+Tree Concurrency Control: Example



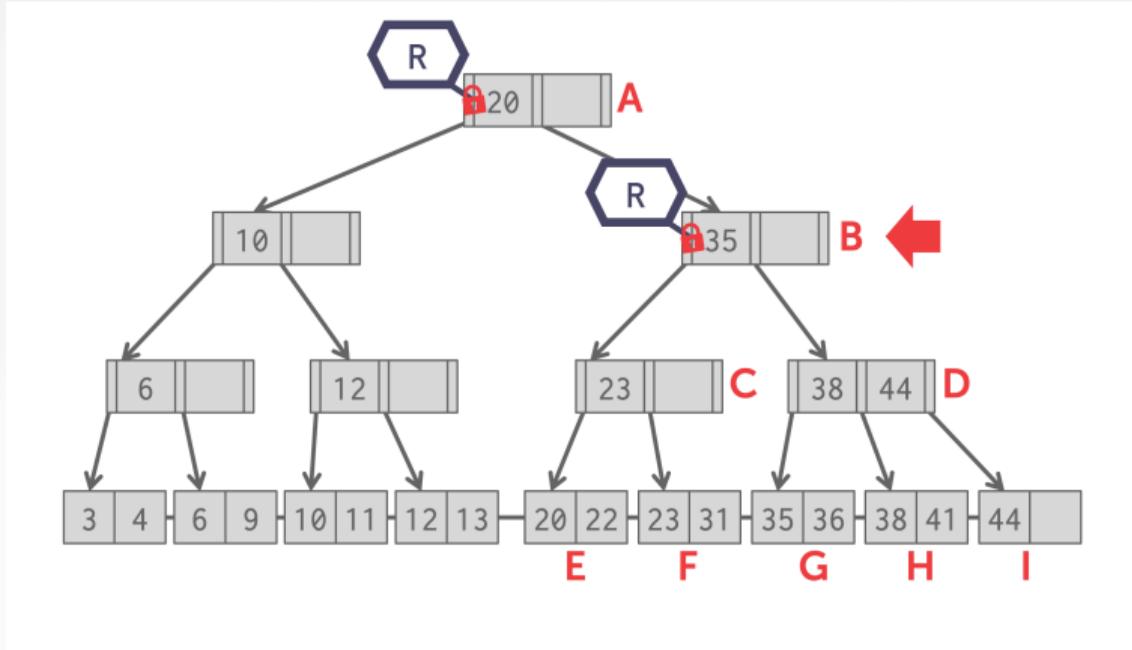




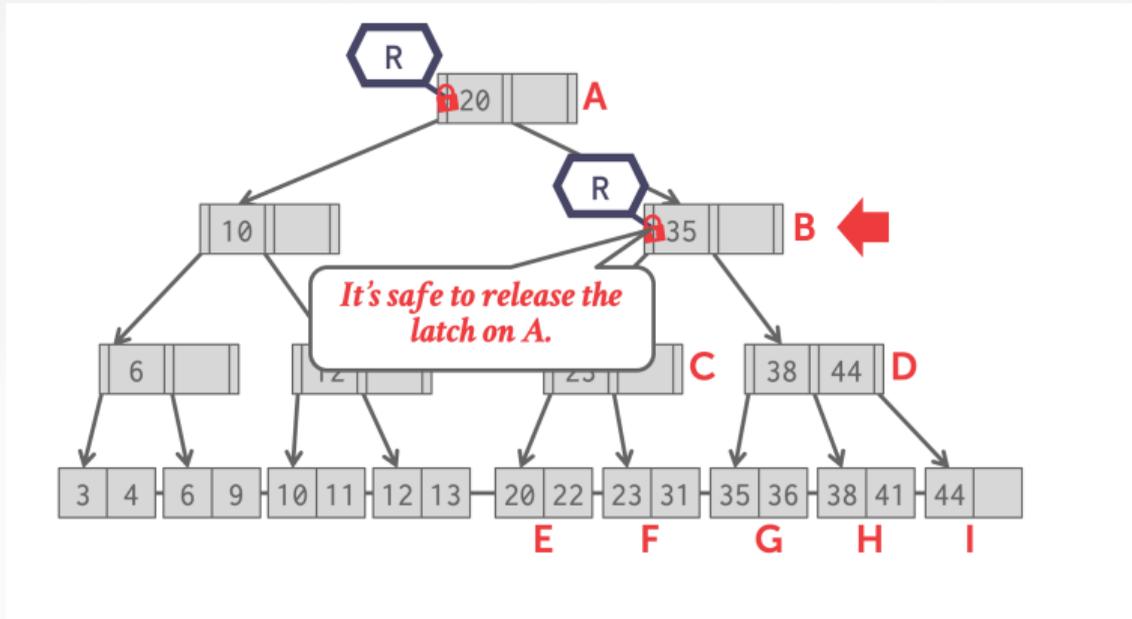
# Example 1 - Find 38



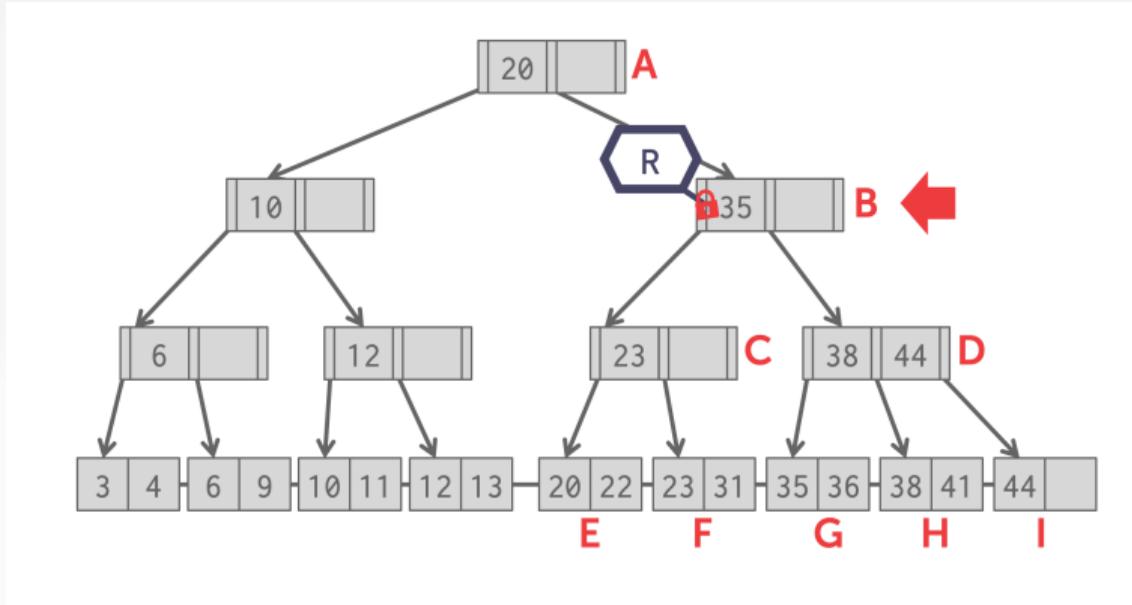
# Example 1 - Find 38



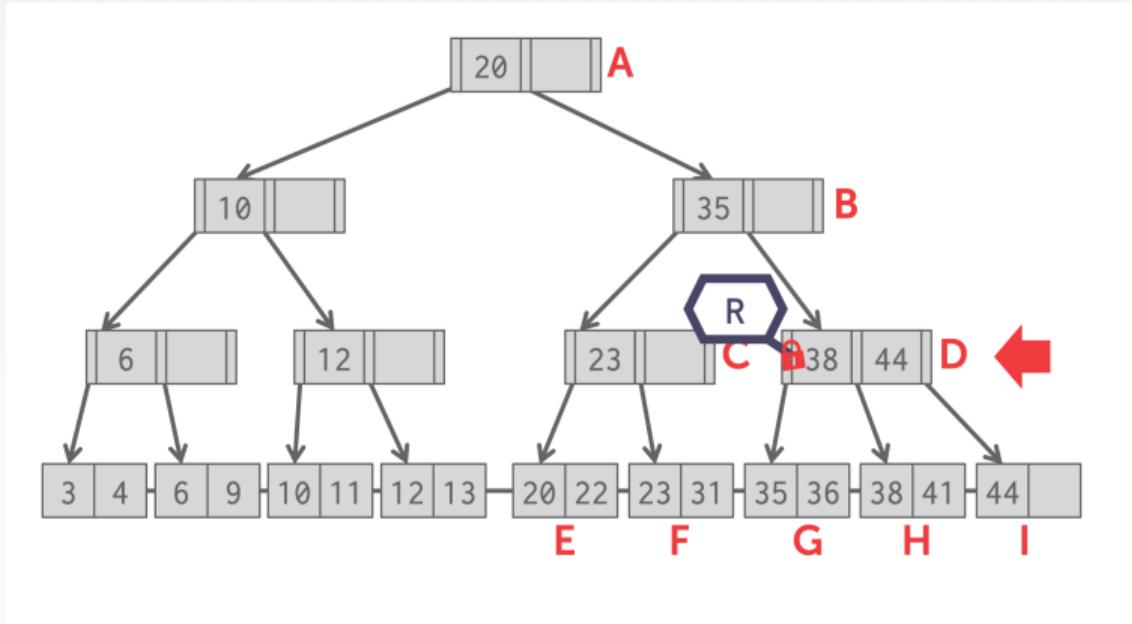
# Example 1 - Find 38



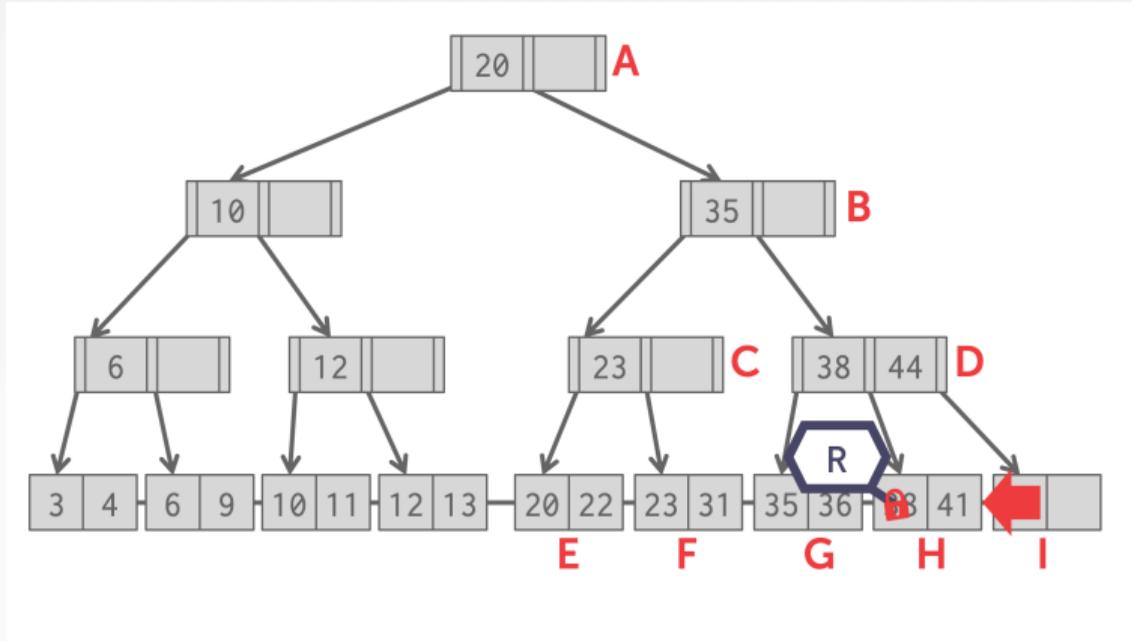
# Example 1 - Find 38



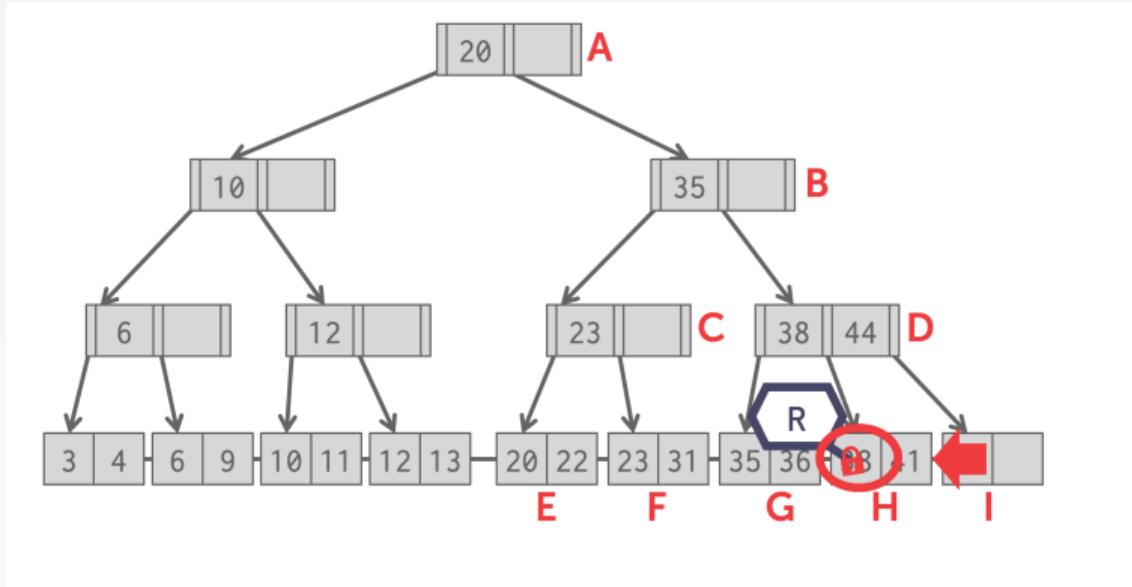
# Example 1 - Find 38



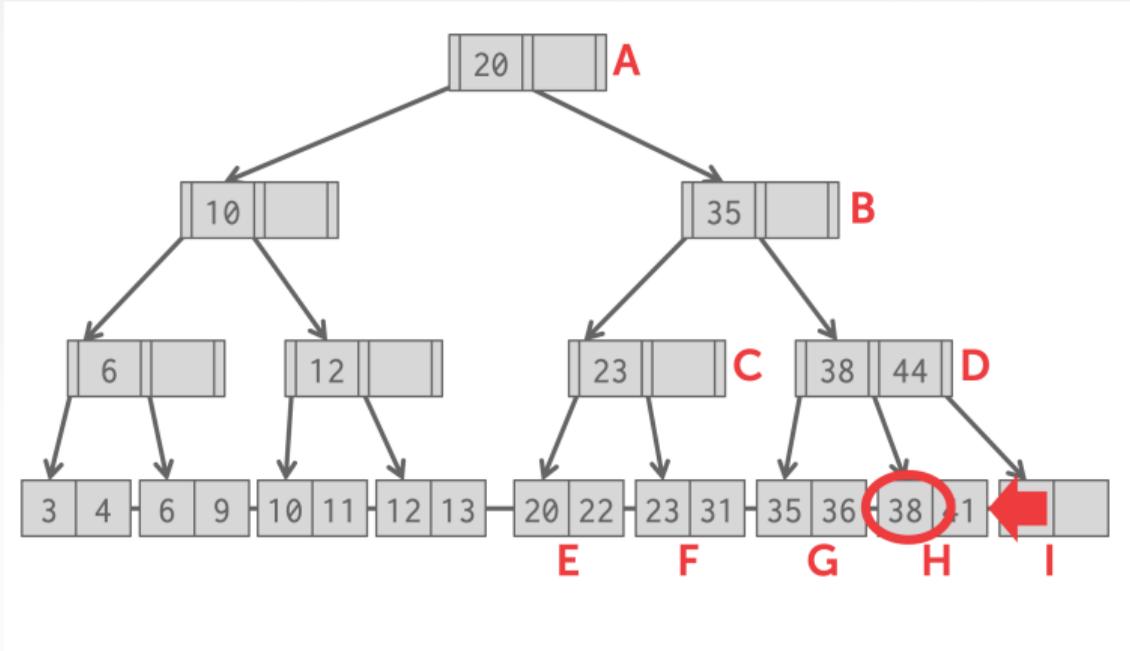
# Example 1 - Find 38



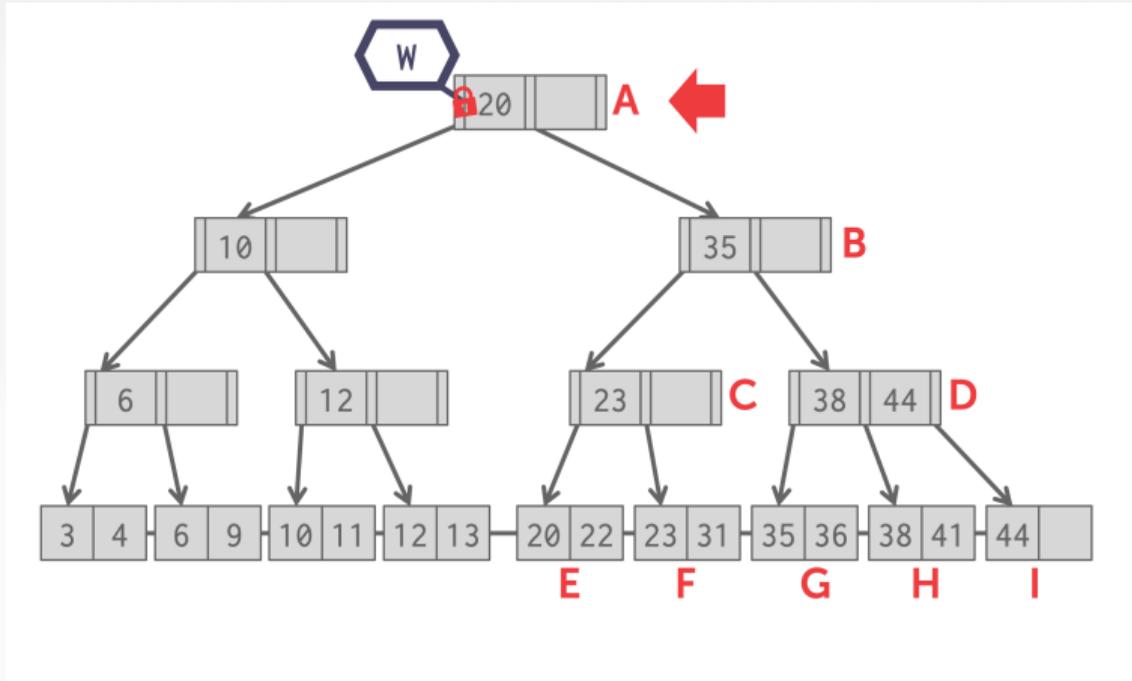
# Example 1 - Find 38



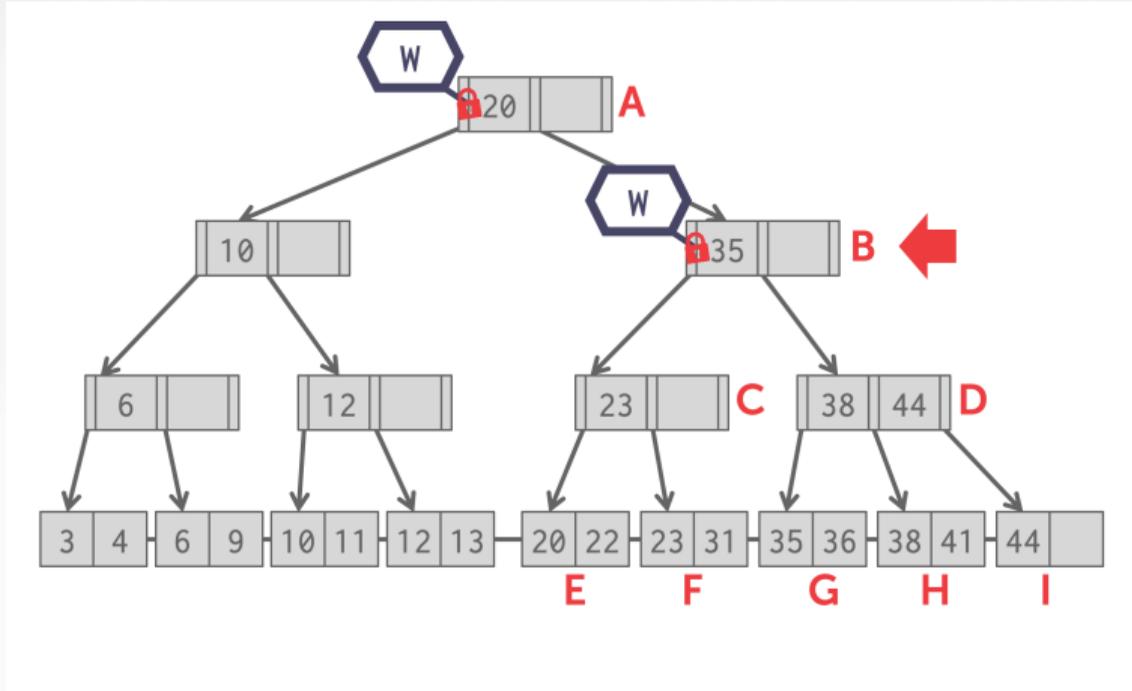
# Example 1 - Find 38



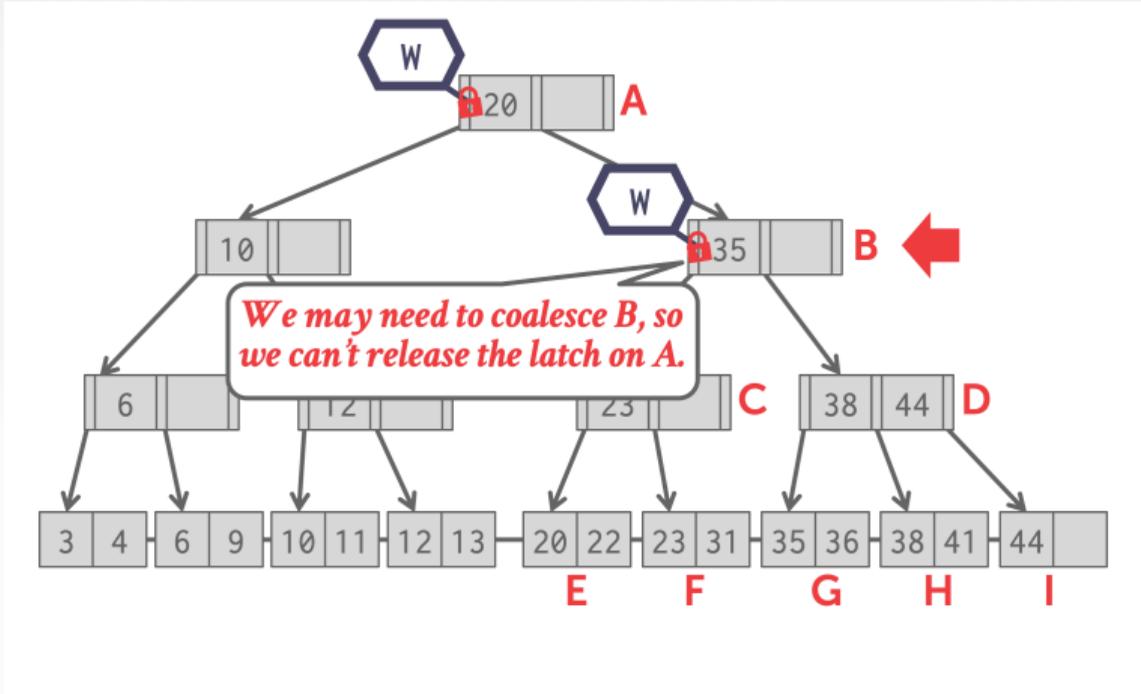
# Example 2 - Delete 38



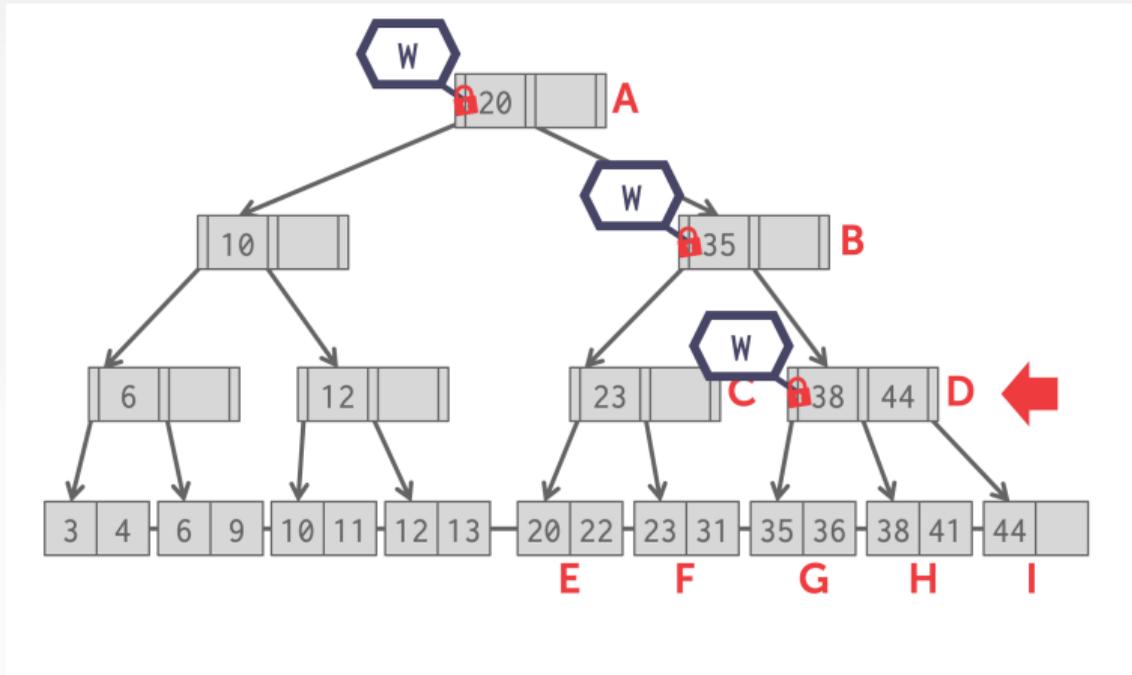
# Example 2 - Delete 38



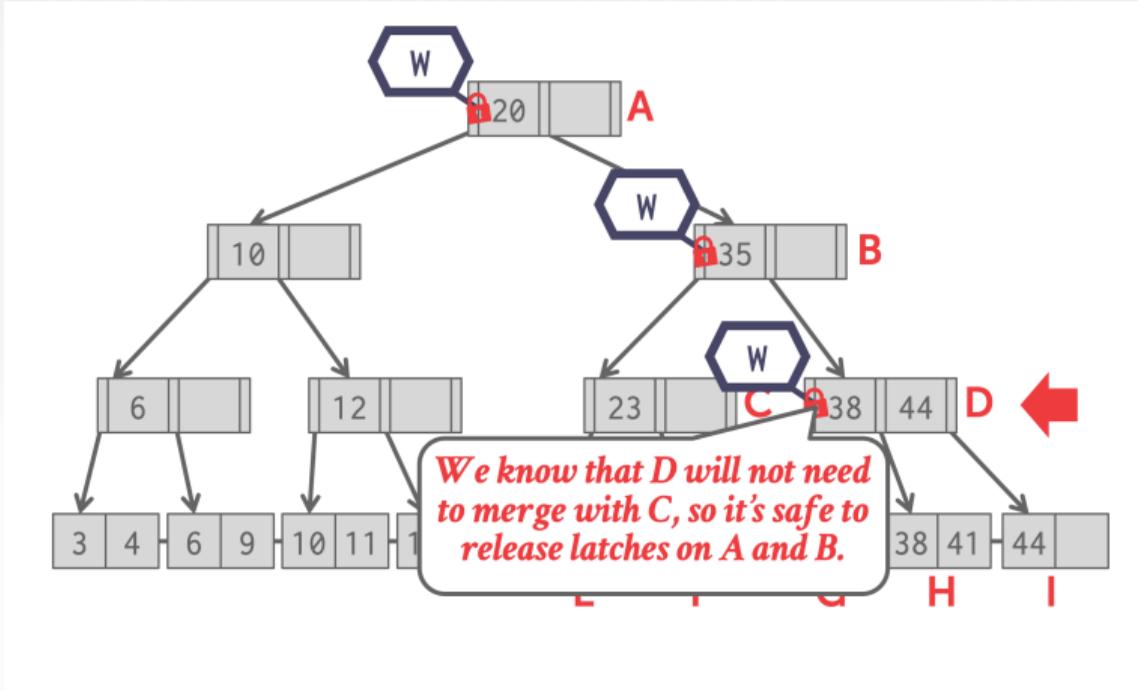
# Example 2 - Delete 38



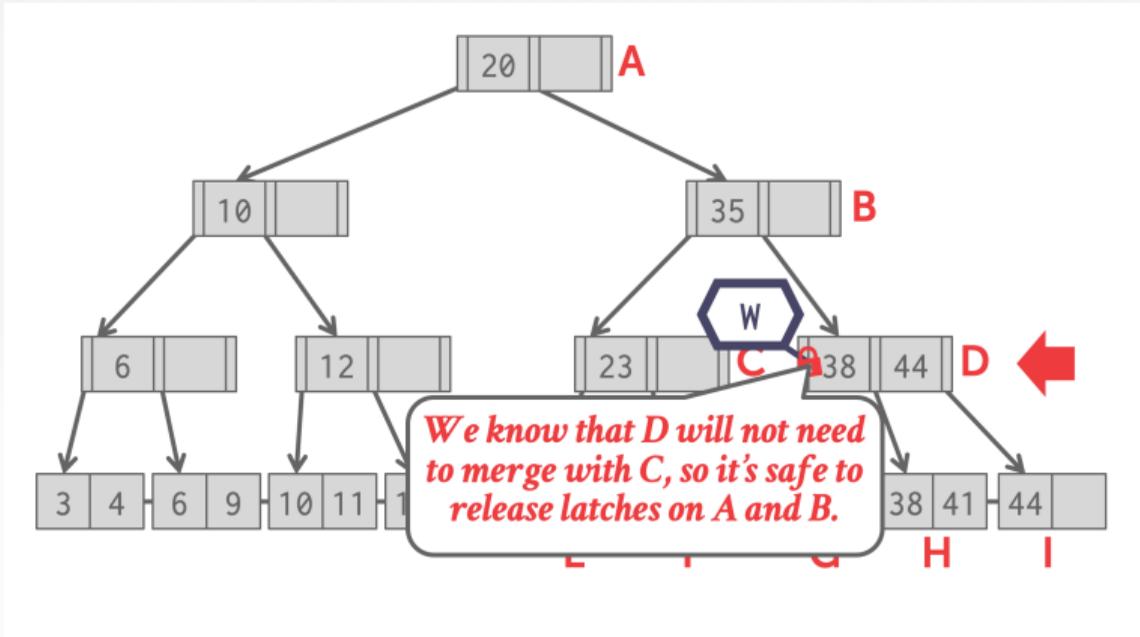
# Example 2 - Delete 38



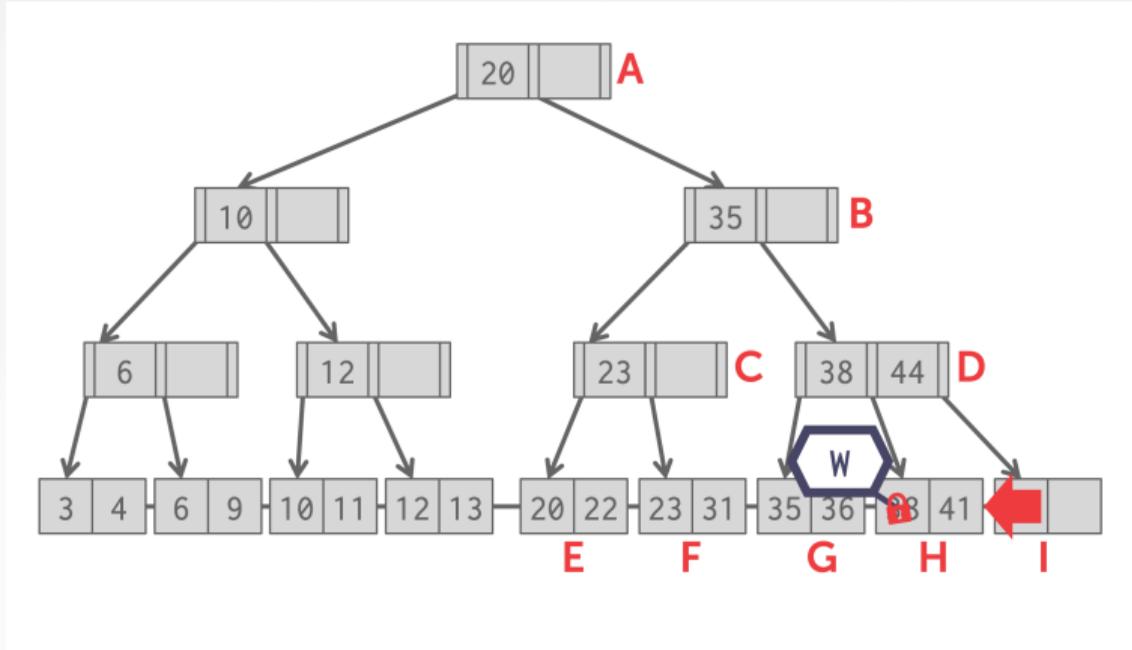
# Example 2 - Delete 38



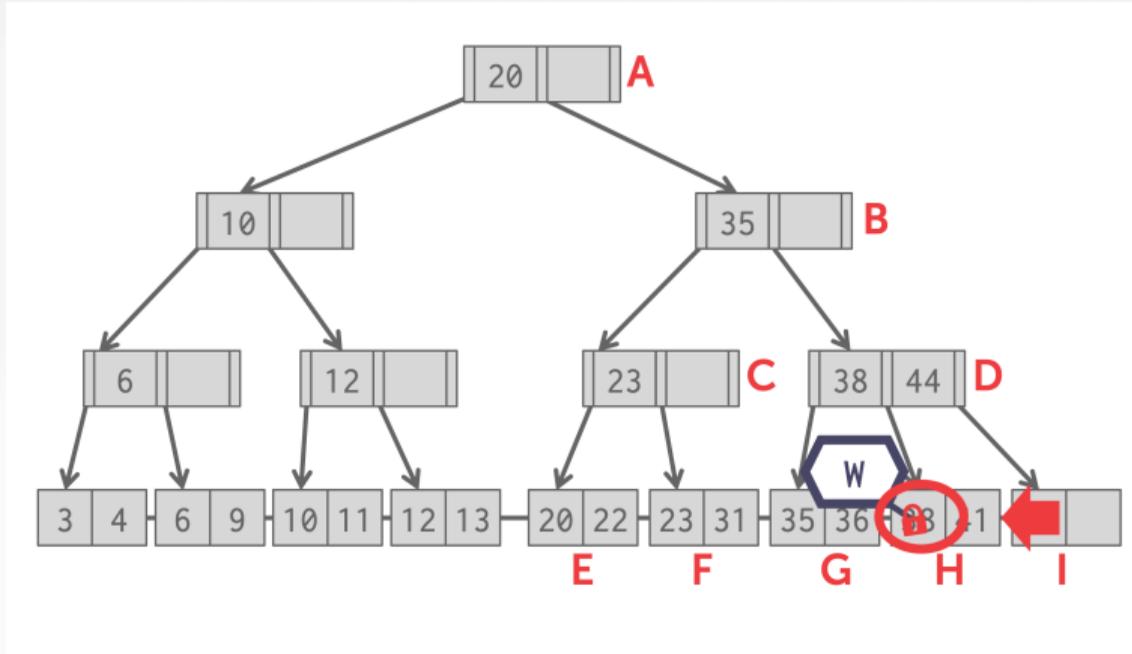
# Example 2 - Delete 38



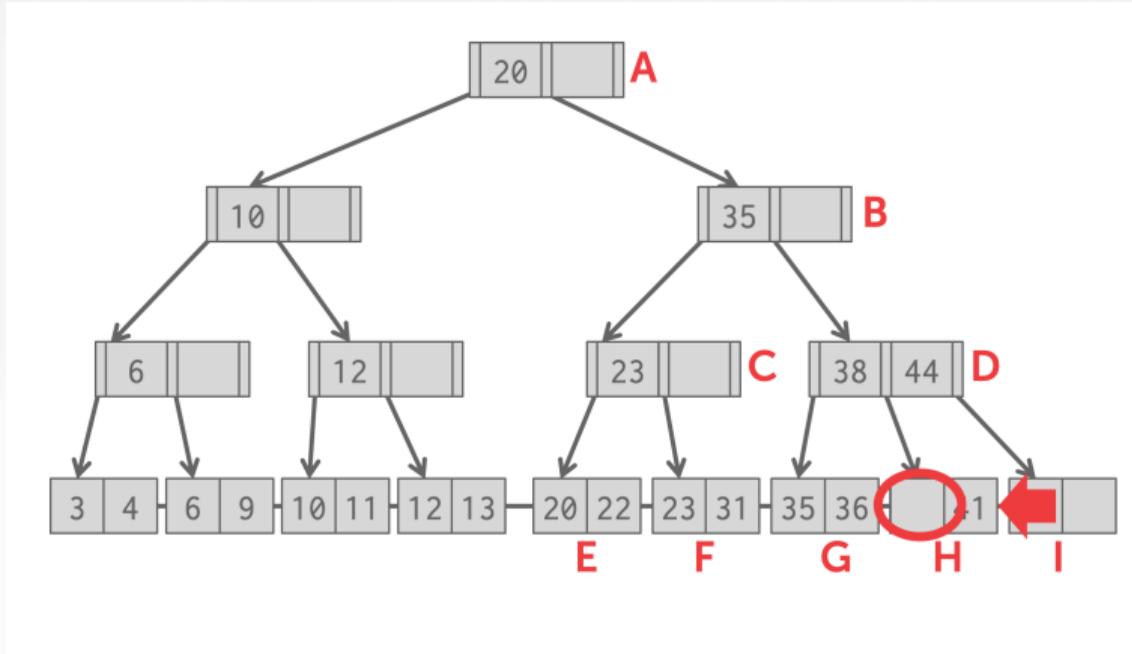
# Example 2 - Delete 38



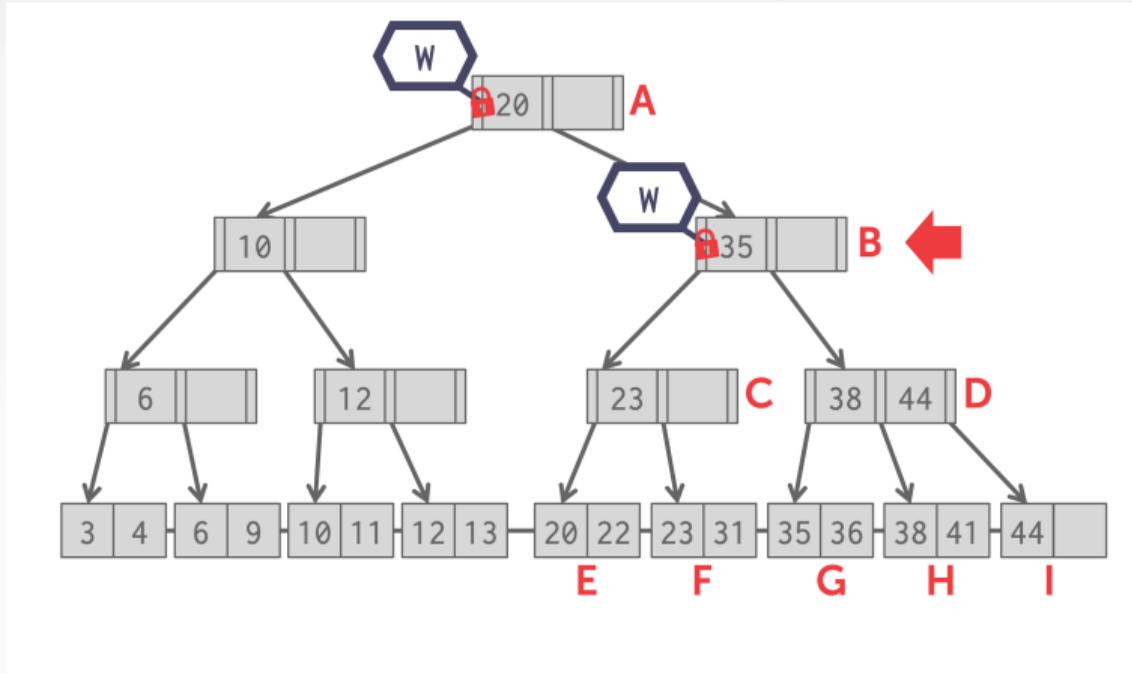
# Example 2 - Delete 38



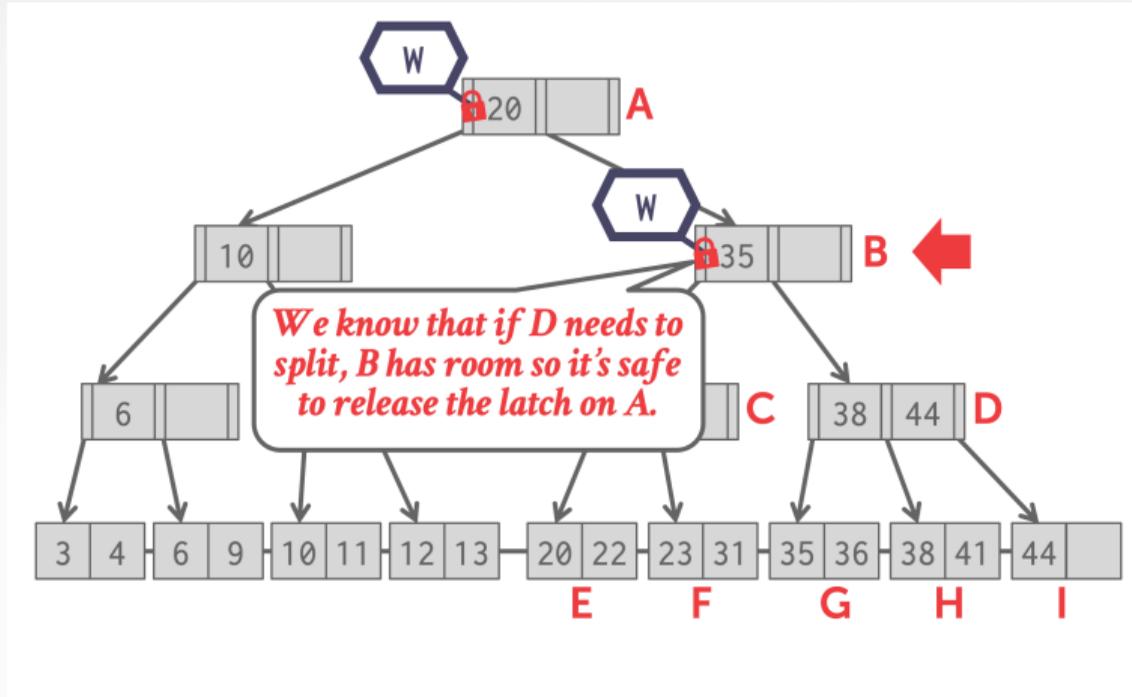
# Example 2 - Delete 38



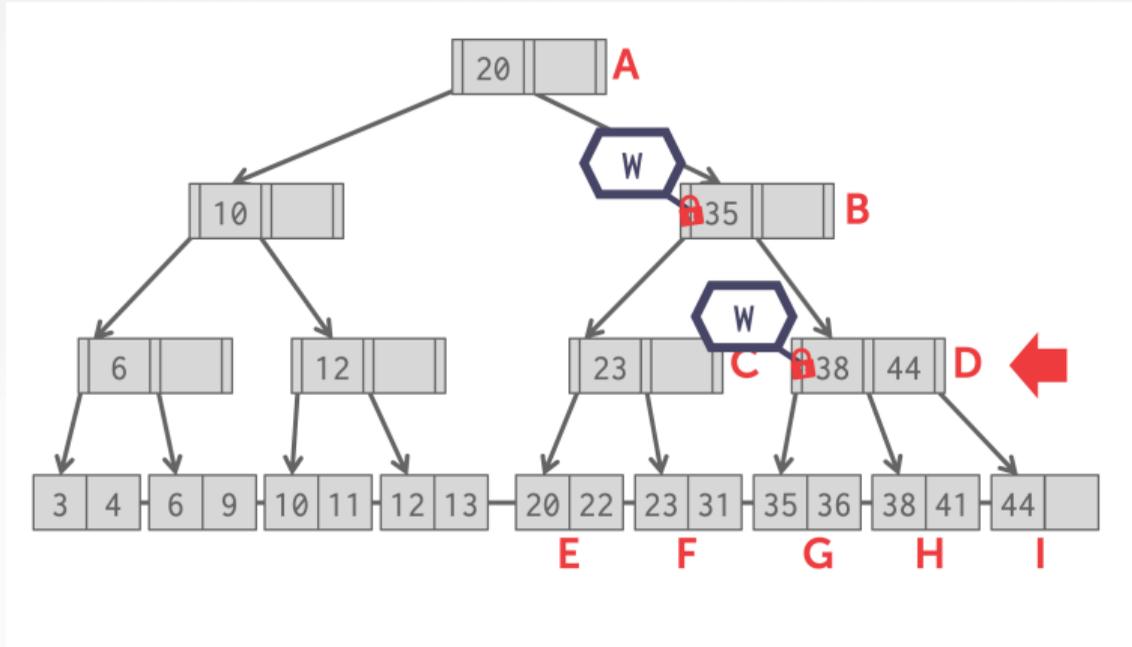
# Example 3 - Insert 45



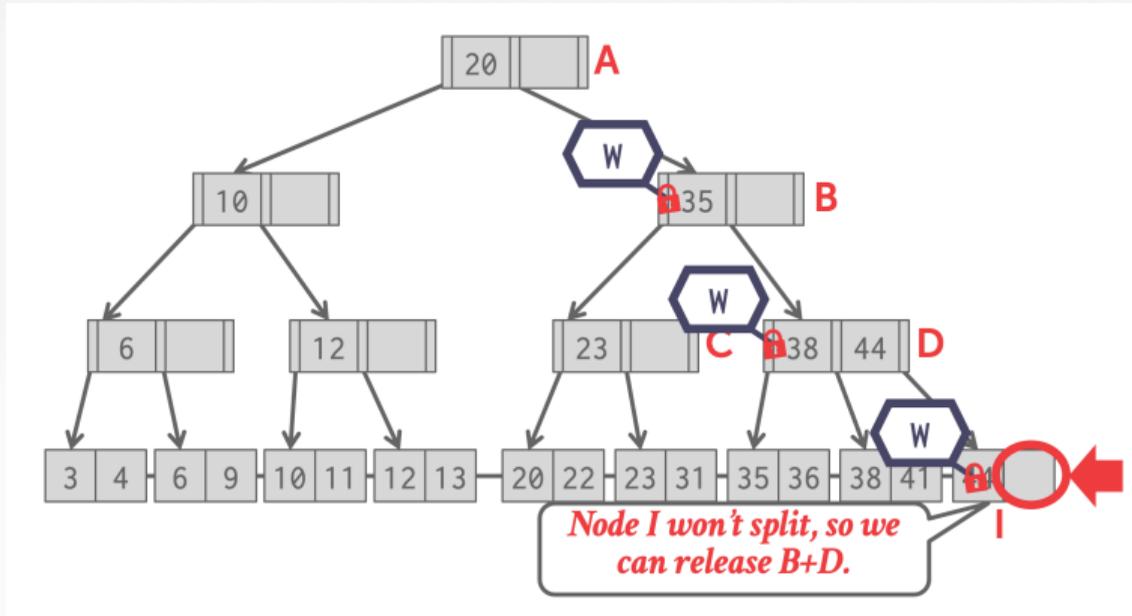
# Example 3 - Insert 45



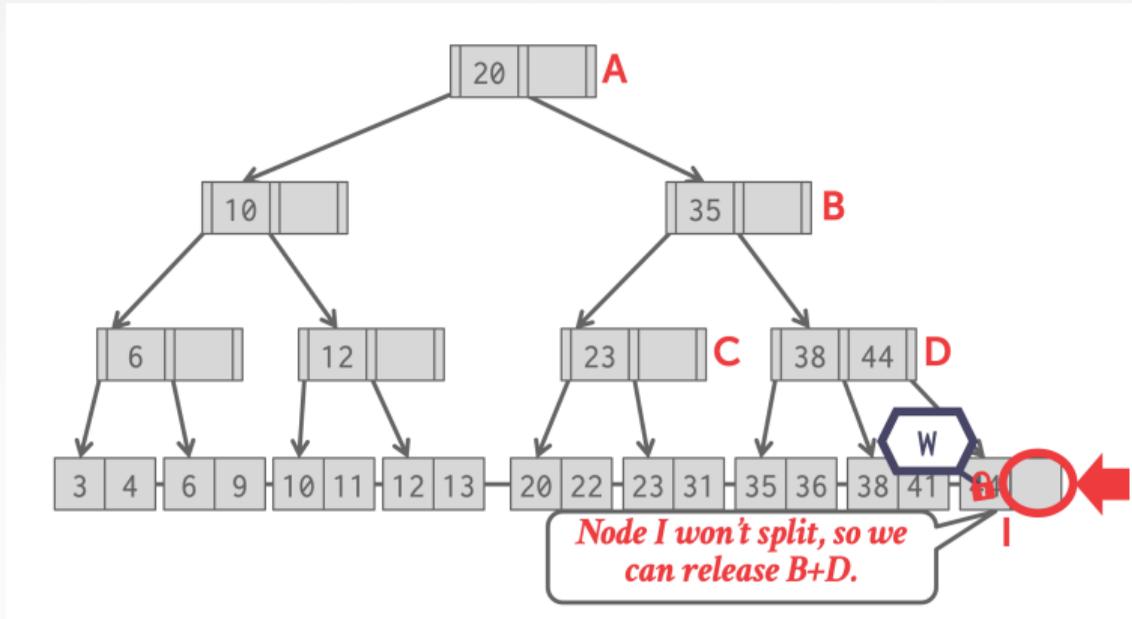
# Example 3 - Insert 45



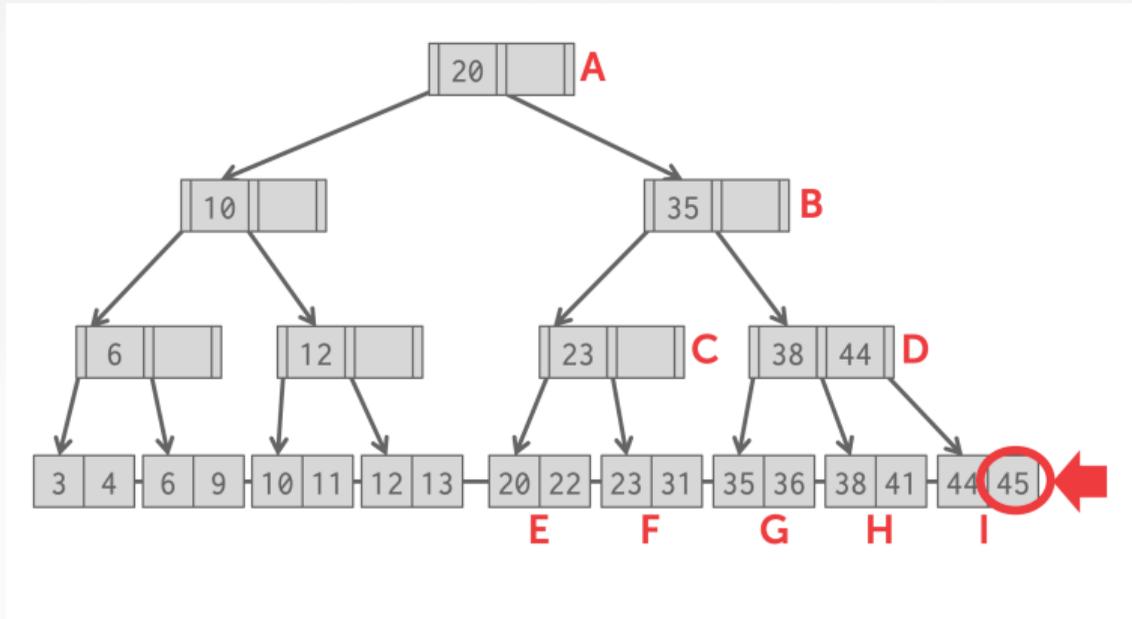
# Example 3 - Insert 45



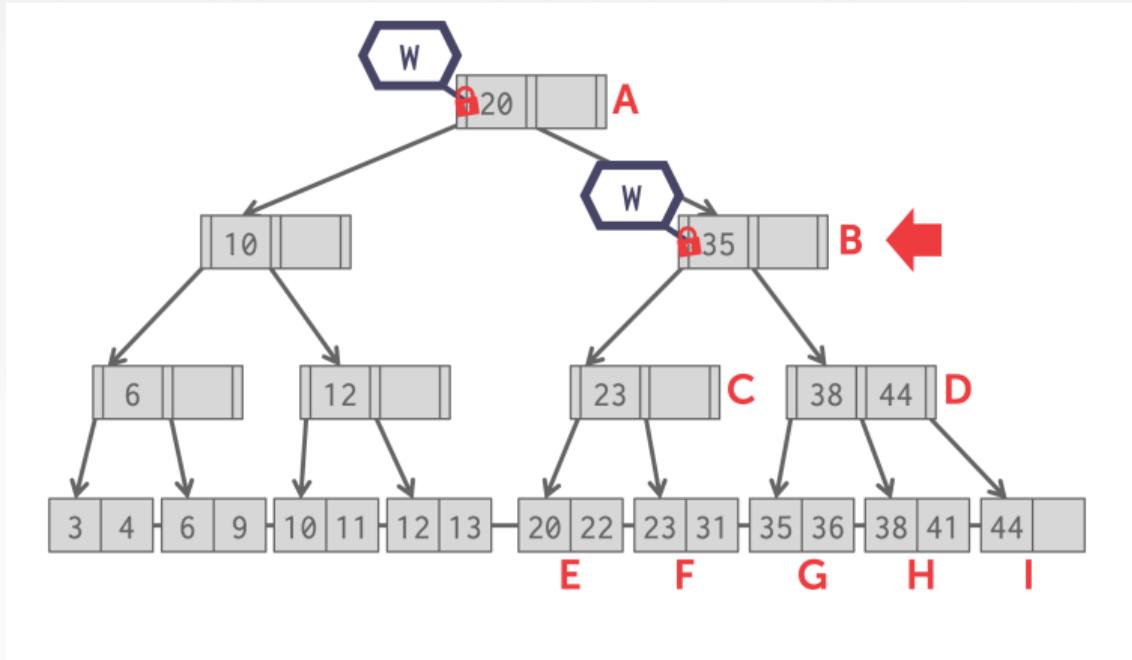
# Example 3 - Insert 45



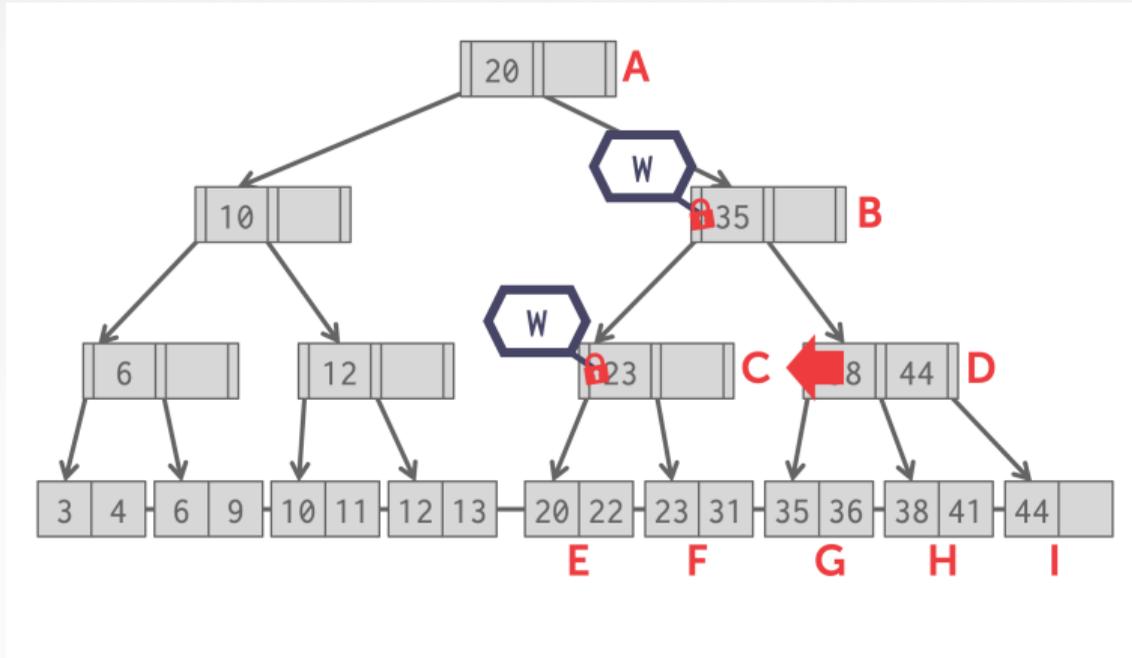
# Example 3 - Insert 45



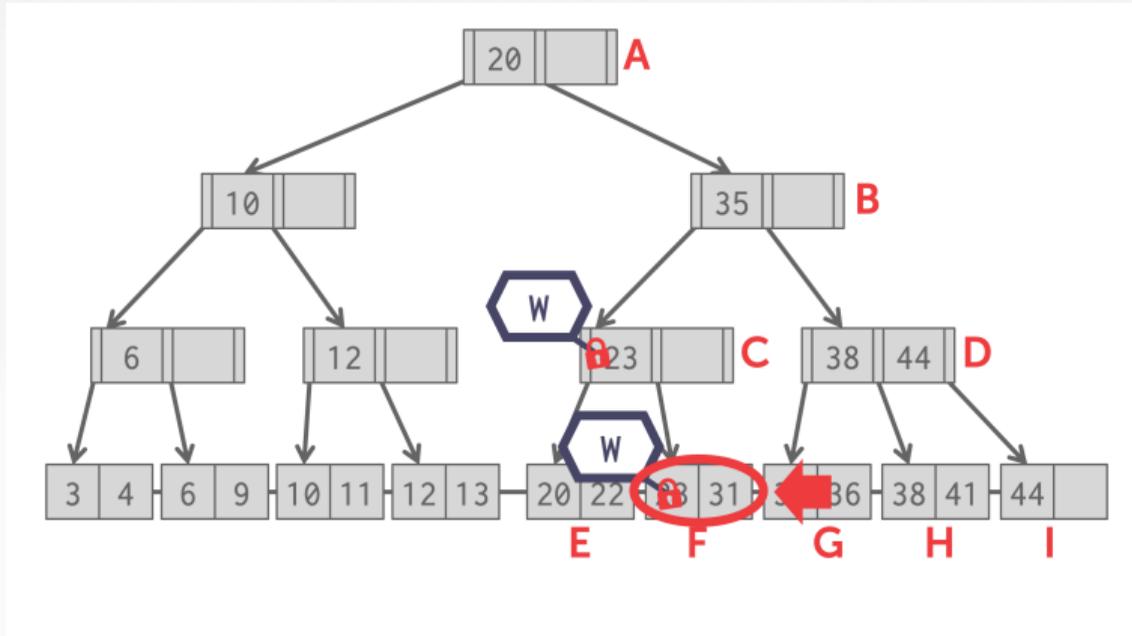
# Example 4 - Insert 25



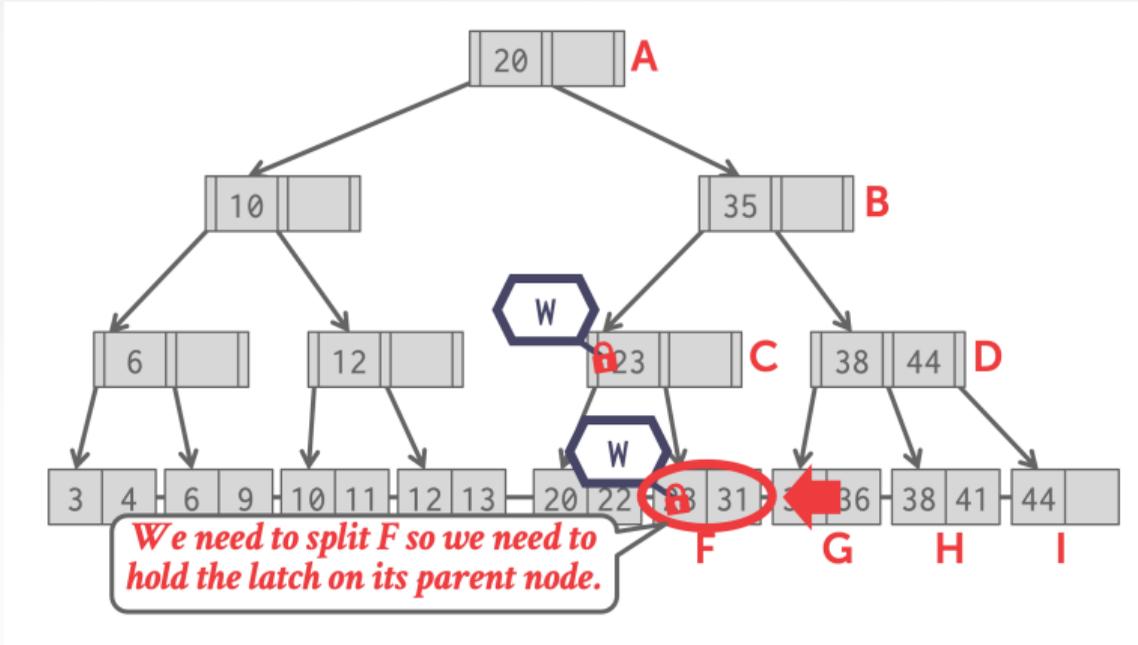
# Example 4 - Insert 25



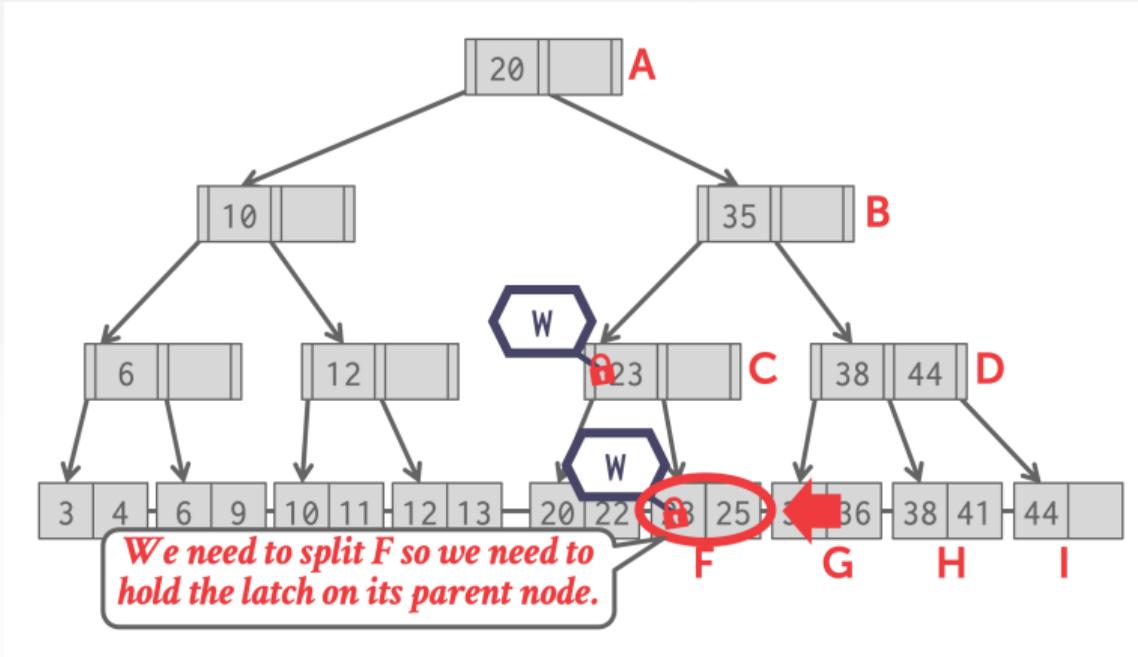
# Example 4 - Insert 25



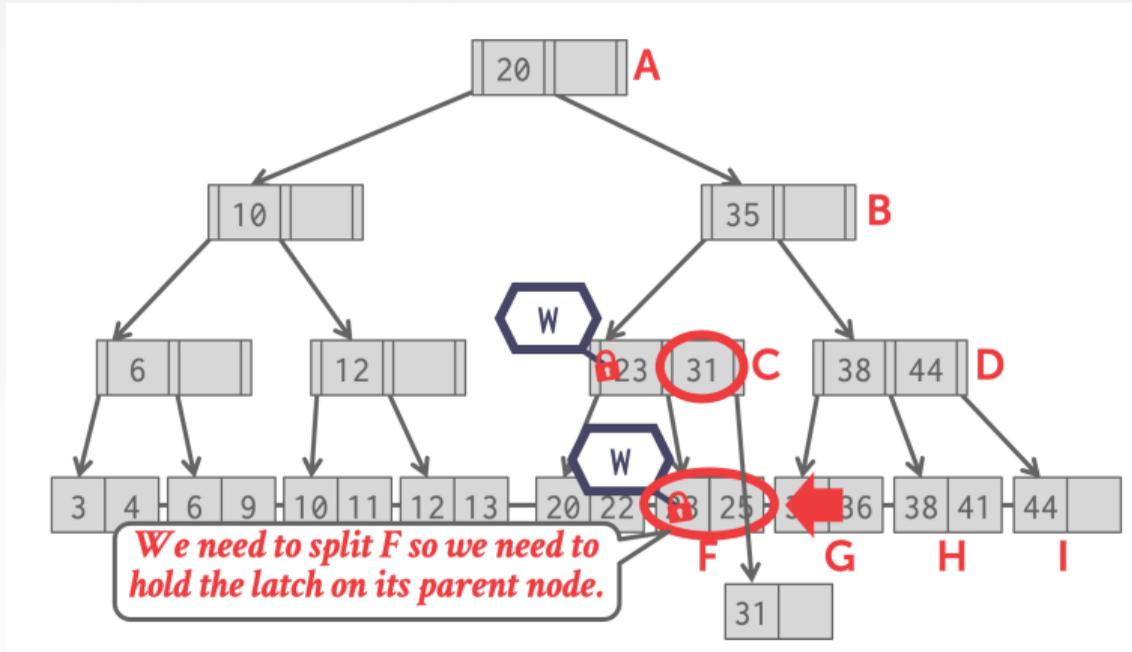
# Example 4 - Insert 25



# Example 4 - Insert 25



# Example 4 - Insert 25

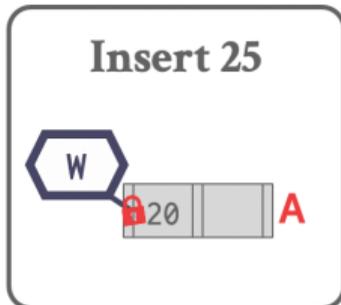
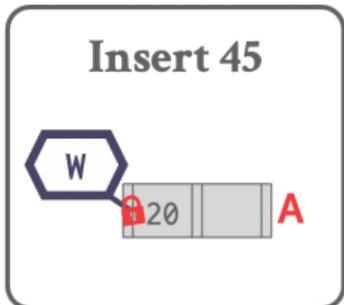
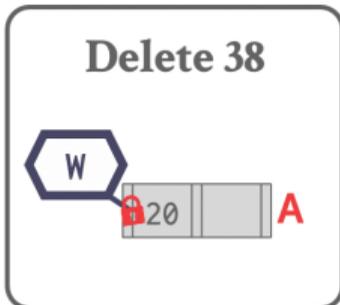




## Observation

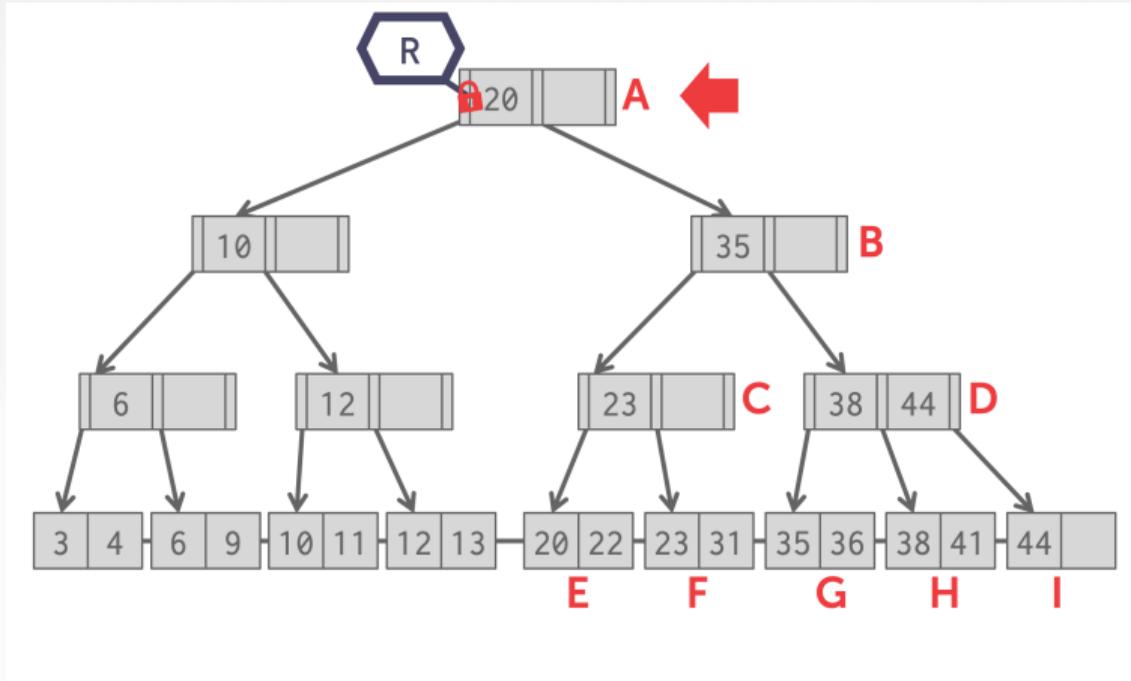
---

- What was the first step that all the update examples did on the B+Tree?
- Taking a write latch on the root every time becomes a bottleneck with higher concurrency.
- Can we do better?

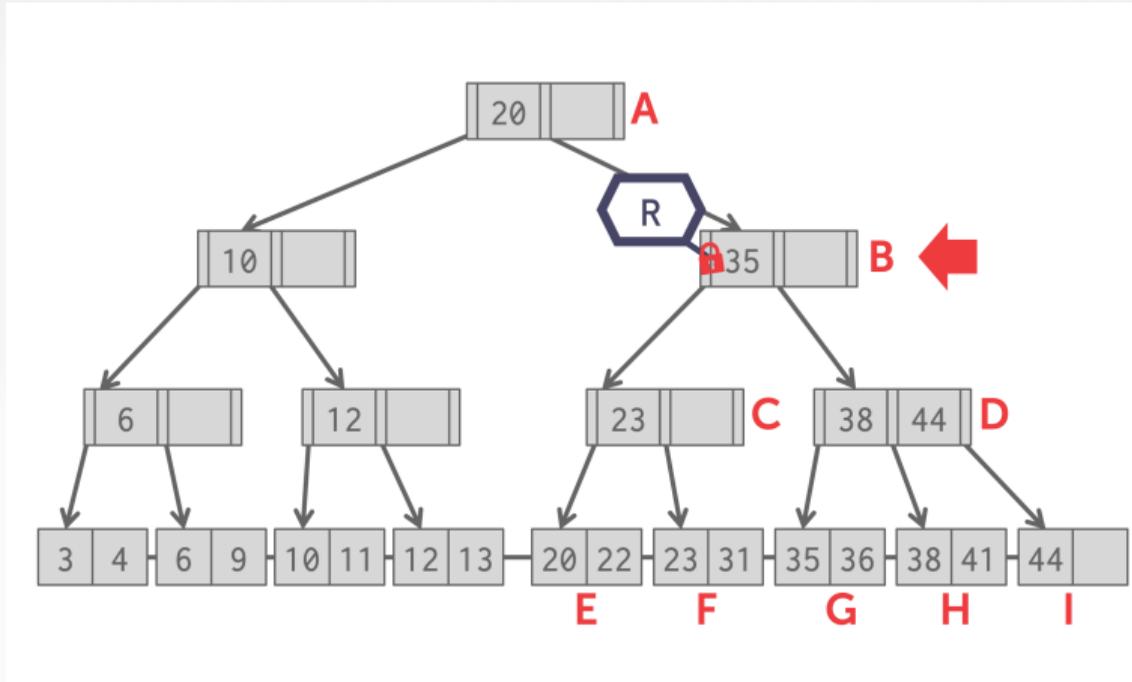




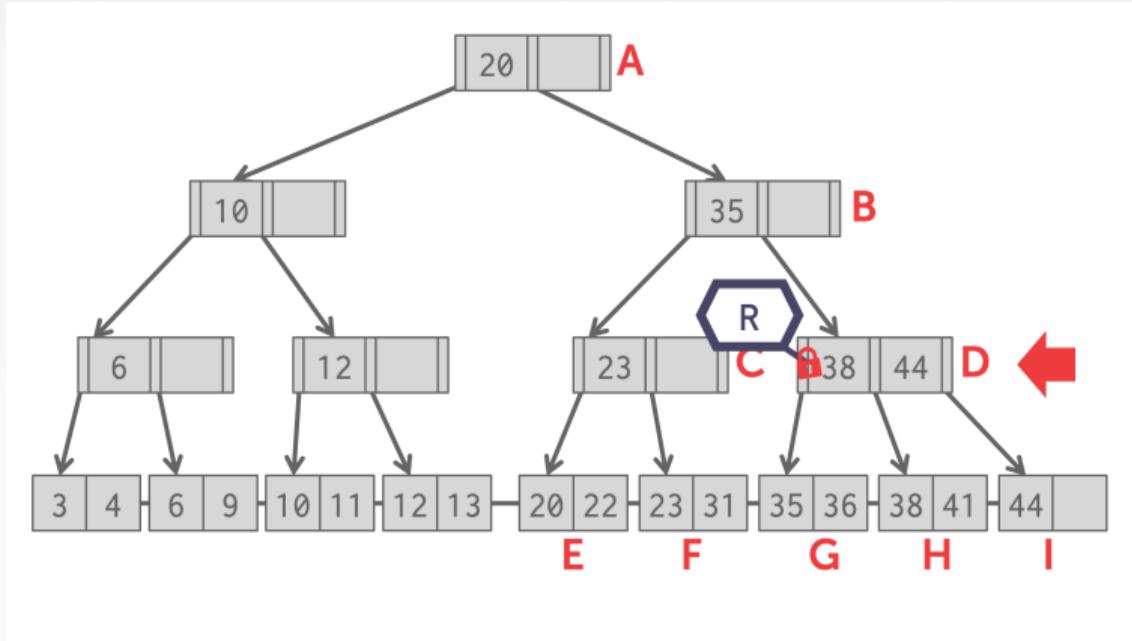
# Example 2 - Delete 38



# Example 2 - Delete 38

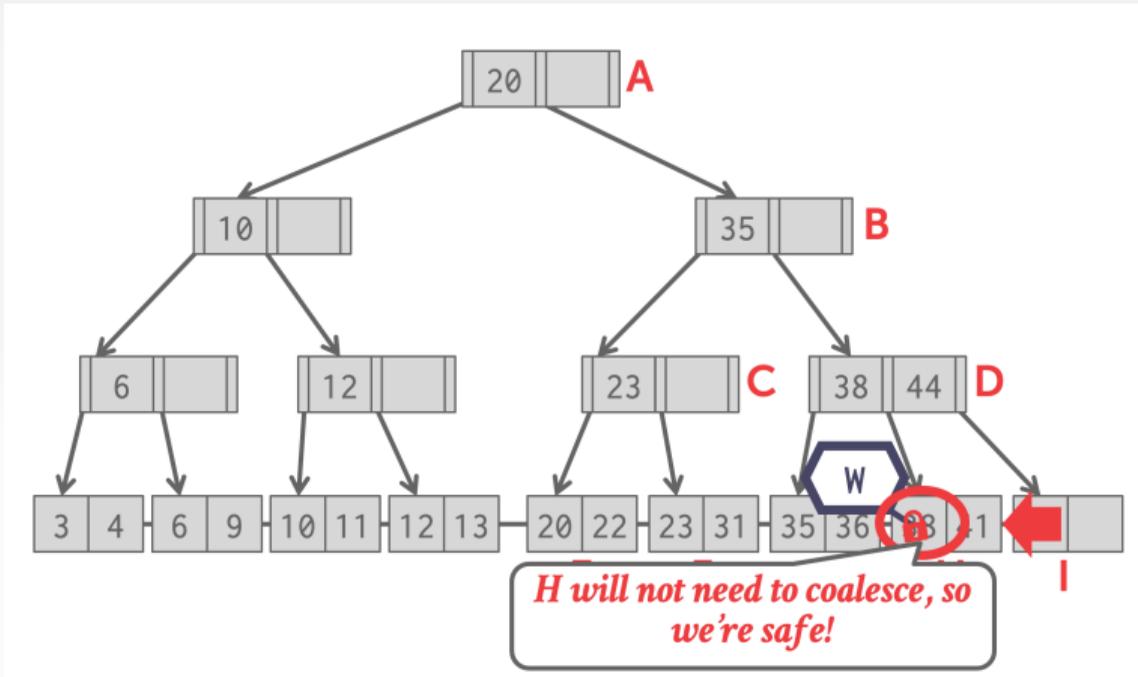


# Example 2 - Delete 38

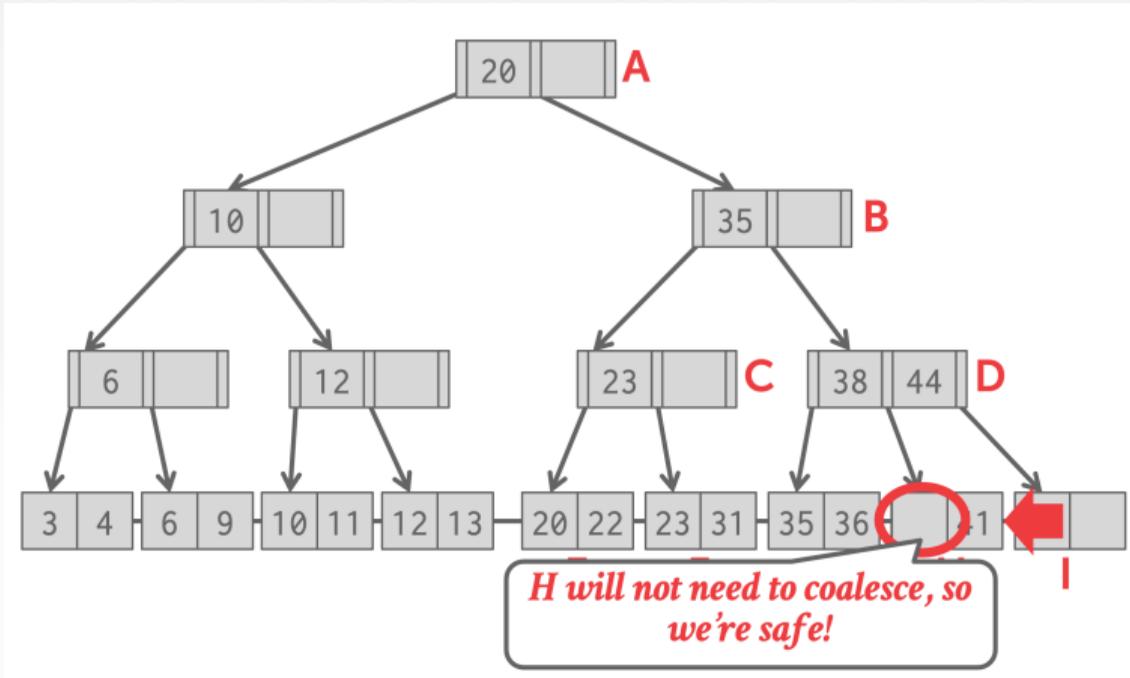




# Example 2 - Delete 38

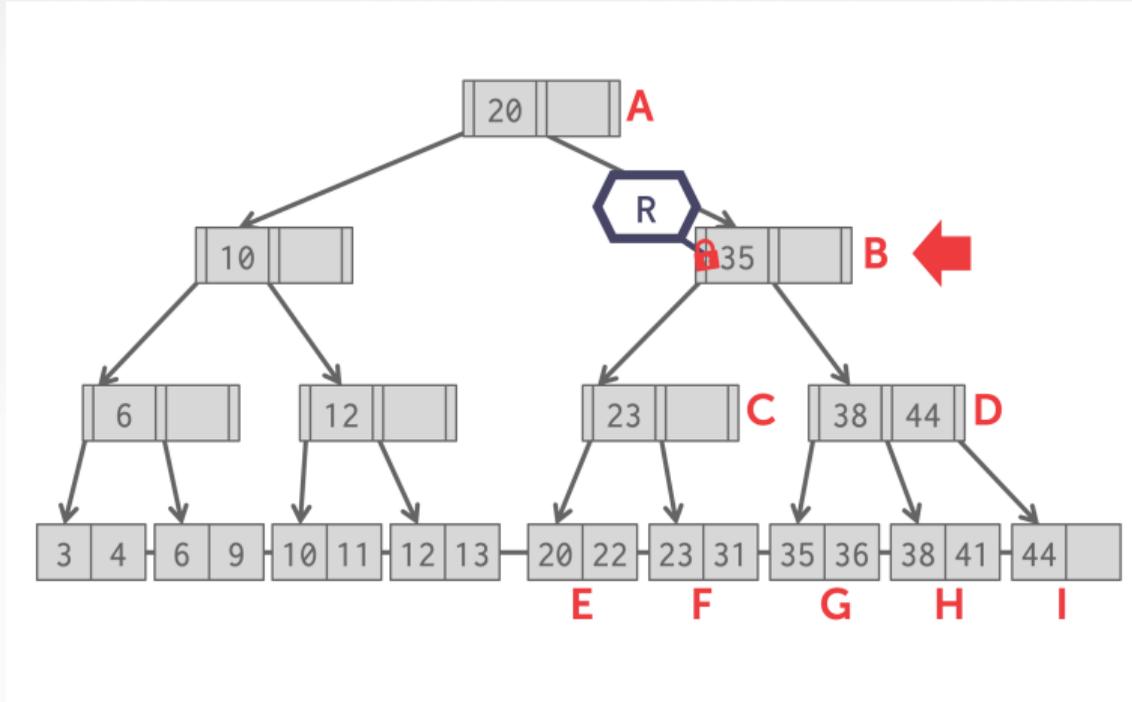


# Example 4 - Insert 25

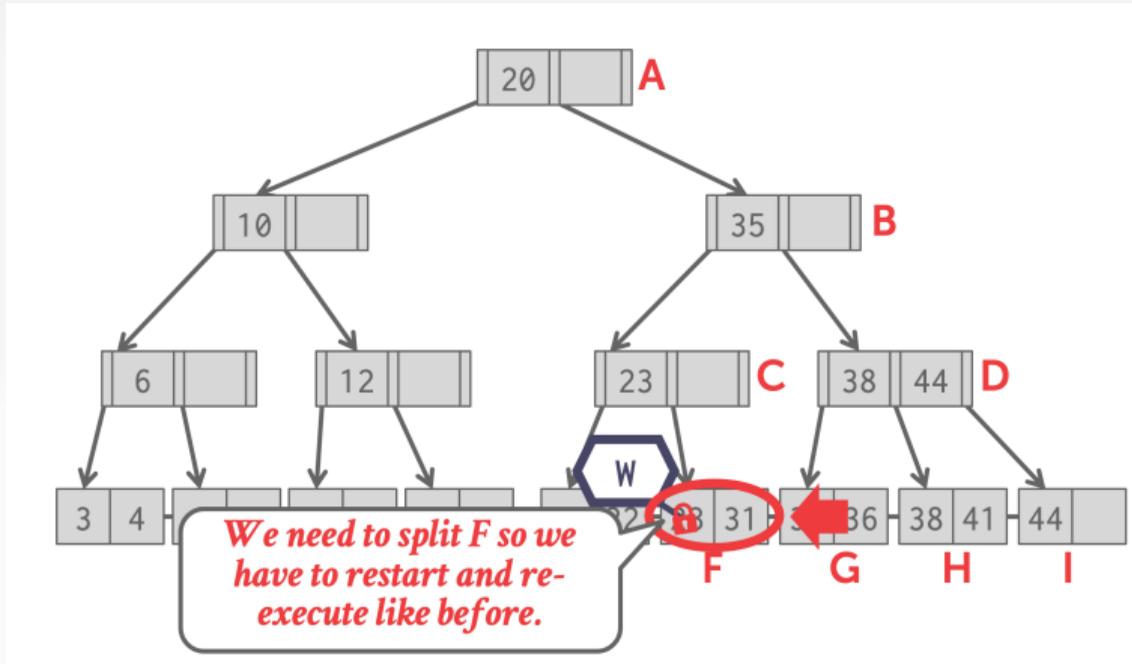




# Example 4 - Insert 25



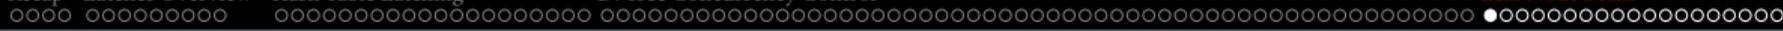
# Example 4 - Insert 25



# Better Latching Algorithm

---

- **Find**: Same as before.
- **Insert/Delete**:
  - ▶ Set latches as if for search, get to leaf, and set **W** latch on leaf.
  - ▶ If leaf is not safe, release all latches, and restart thread using previous insert/delete protocol with **W** latches.
- This approach **optimistically** assumes that only leaf node will be modified; if not, **R** latches set on the first pass to leaf are wasteful.

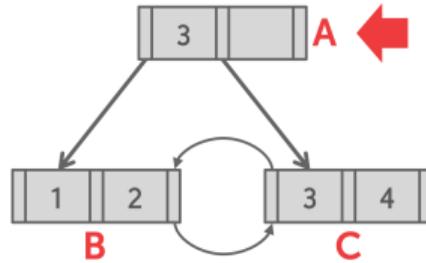


# Leaf Node Scans



# Leaf Node Scan - Example 1

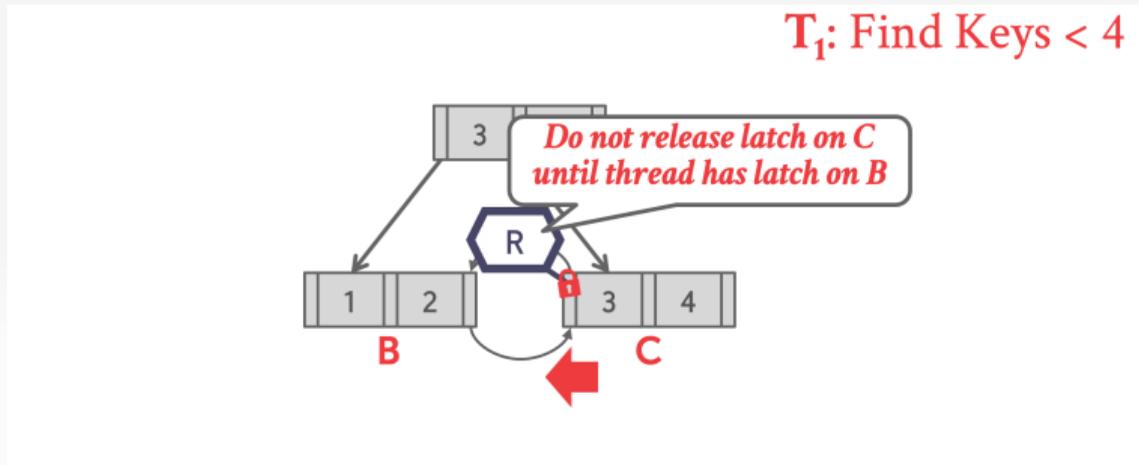
$T_1$ : Find Keys < 4



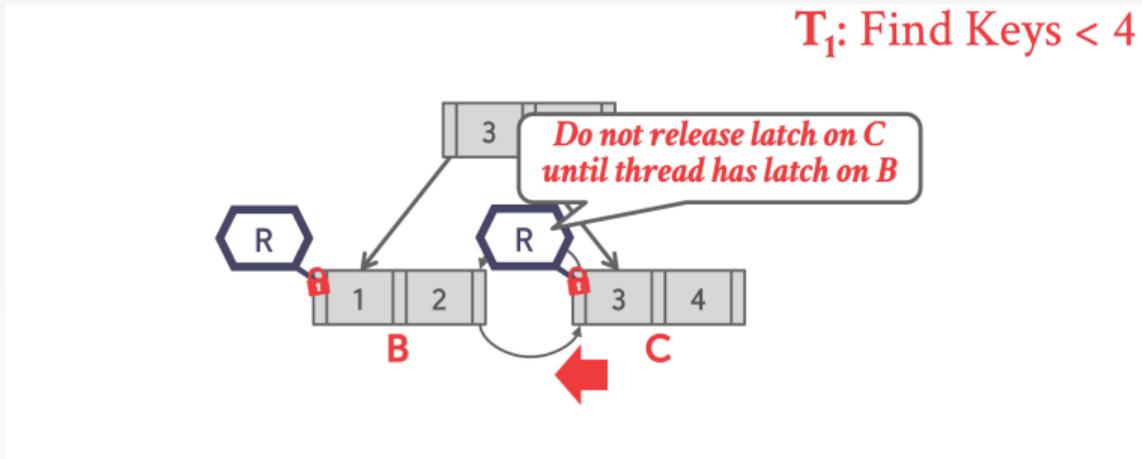




# Leaf Node Scan - Example 1

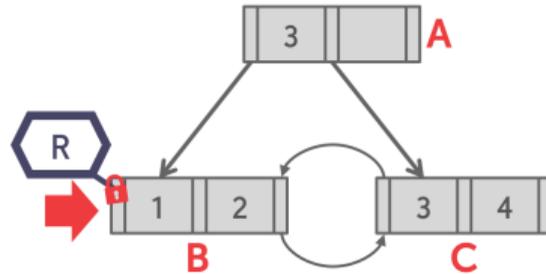


# Leaf Node Scan - Example 1



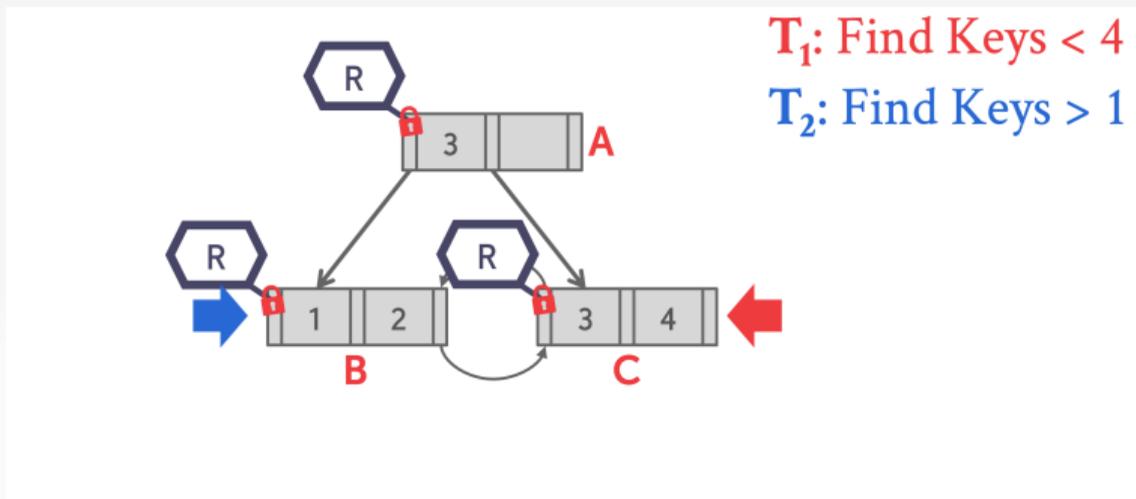
# Leaf Node Scan - Example 1

$T_1$ : Find Keys < 4

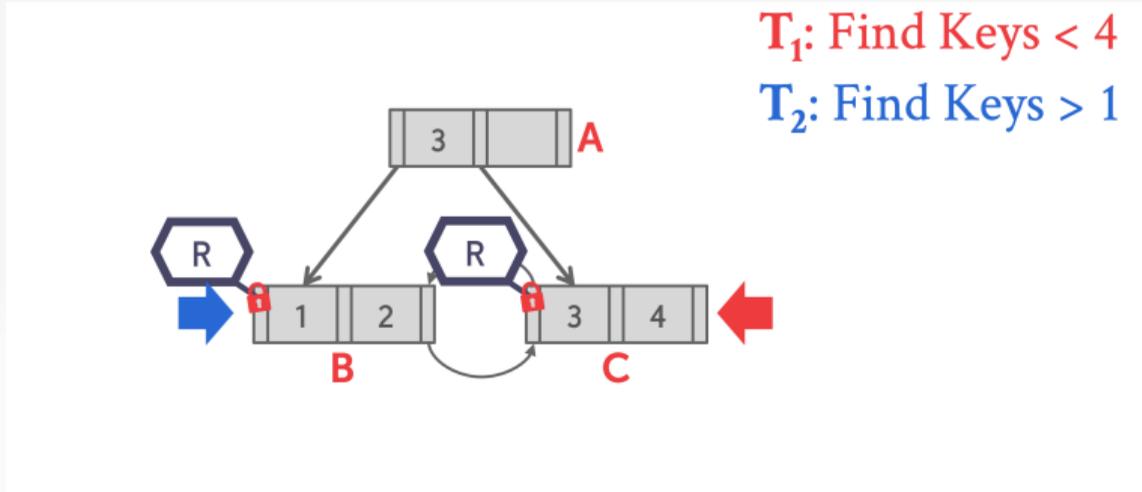




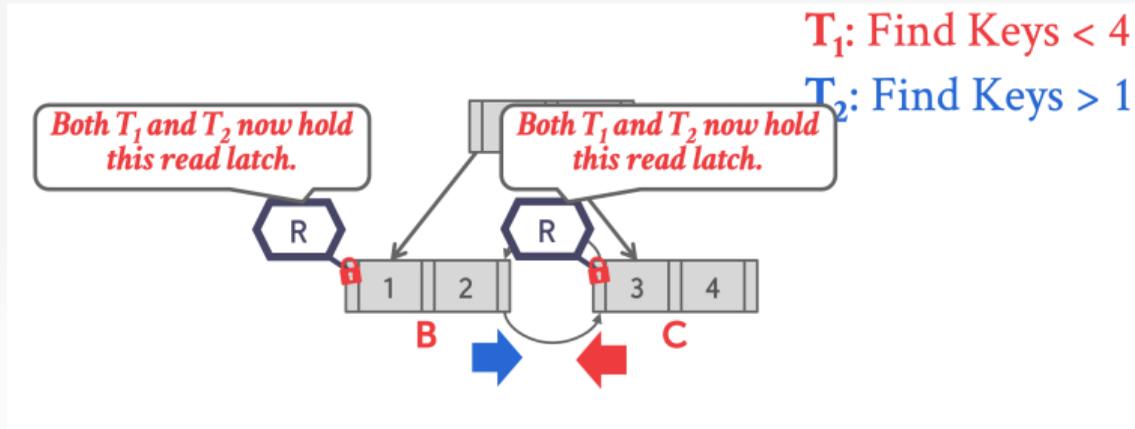
## Leaf Node Scan - Example 2



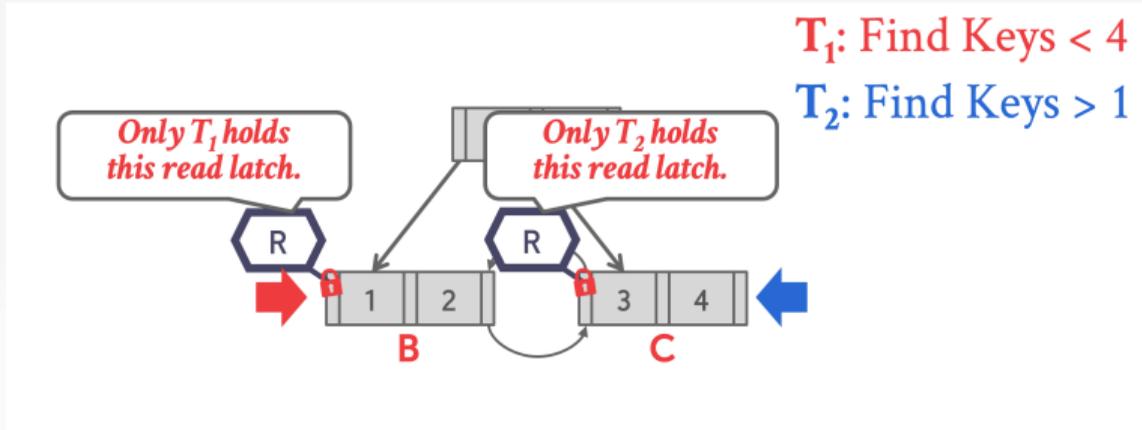
# Leaf Node Scan - Example 2



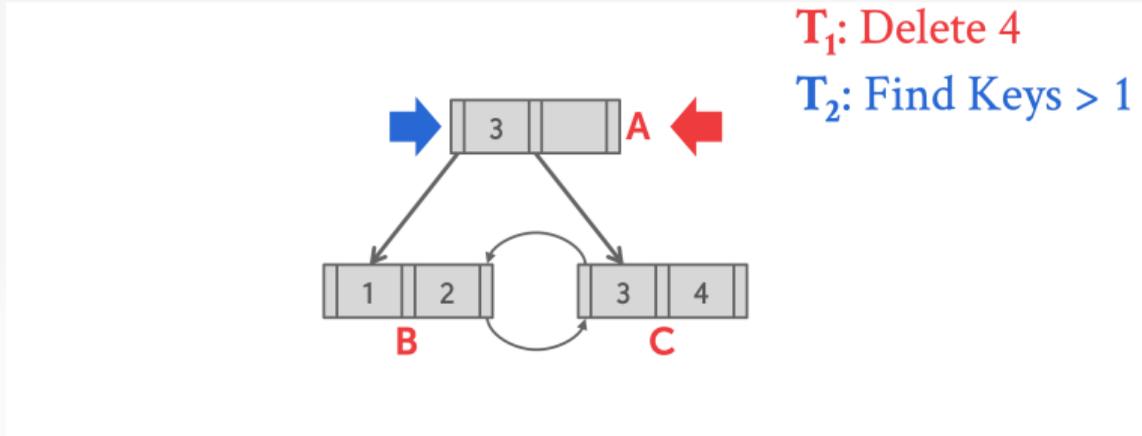
## Leaf Node Scan - Example 2



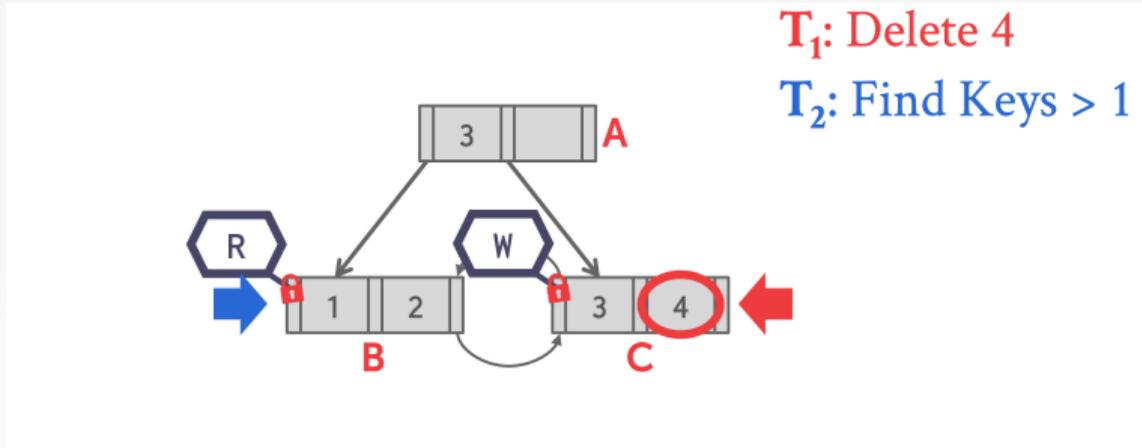
# Leaf Node Scan - Example 2



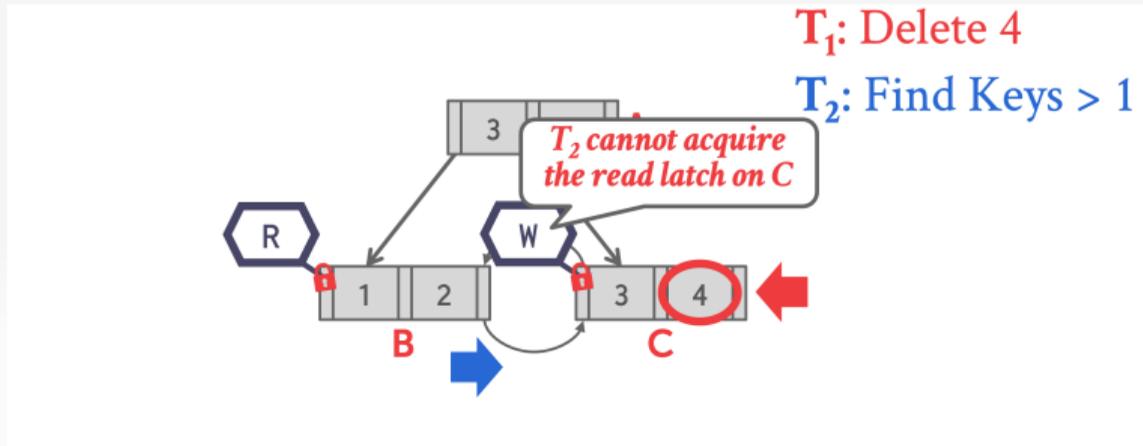
# Leaf Node Scan - Example 3



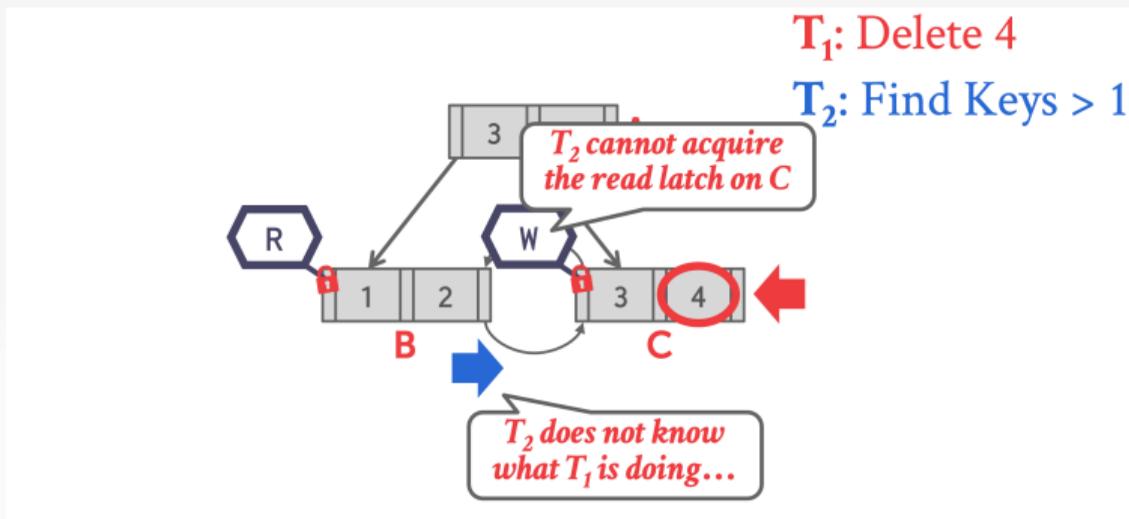
# Leaf Node Scan - Example 3



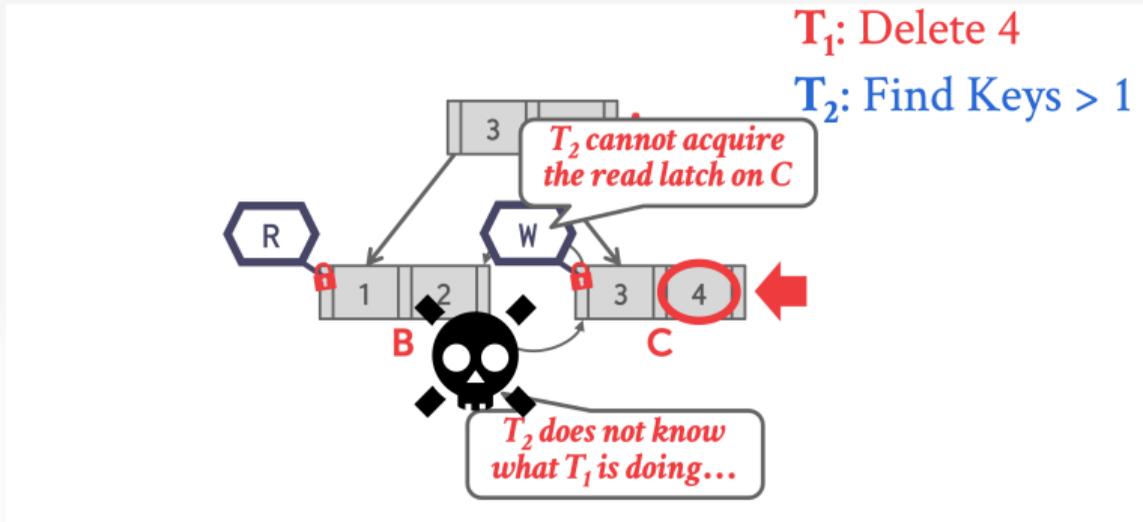
# Leaf Node Scan - Example 3



# Leaf Node Scan - Example 3



# Leaf Node Scan - Example 3



## Leaf Node Scans

---

- Latches do **not** support deadlock detection or avoidance.
- The only way we can deal with this problem is through **coding discipline**.
- The leaf node sibling latch acquisition protocol must support a fail-fast **no-wait** mode.
- B+Tree implementation must cope with failed latch acquisitions.

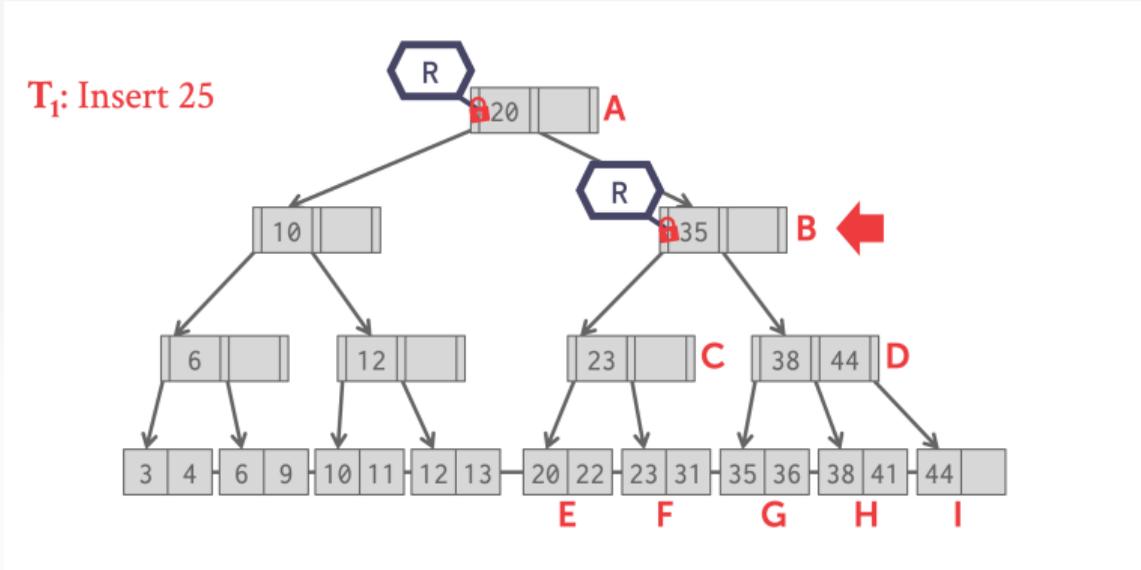


## *B*<sup>link</sup>-Tree

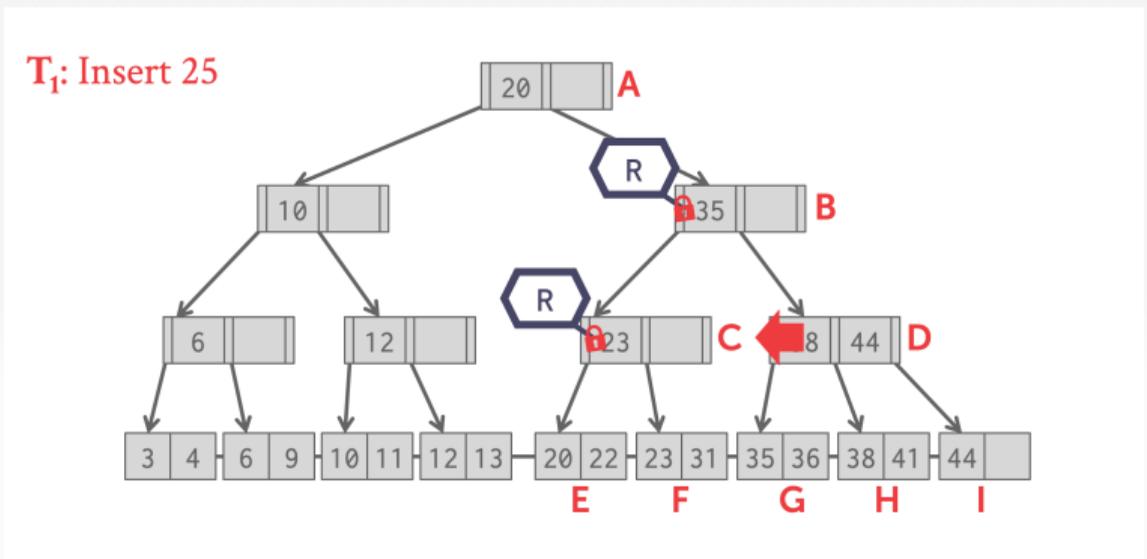
---

- Every time a leaf node overflows, we must update at least **three** nodes.
  - ▶ The leaf node being split.
  - ▶ The new leaf node being created.
  - ▶ The parent node.
- **Optimization:** When a leaf node overflows, delay updating its parent node.
- Reference

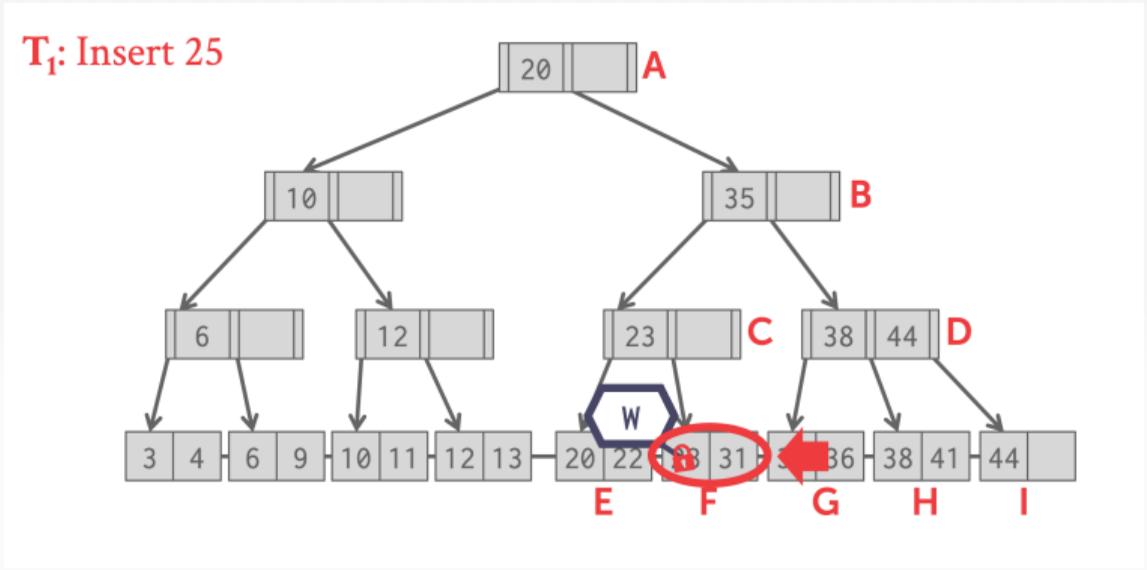
# Blink-Tree Example



# Blink-Tree Example

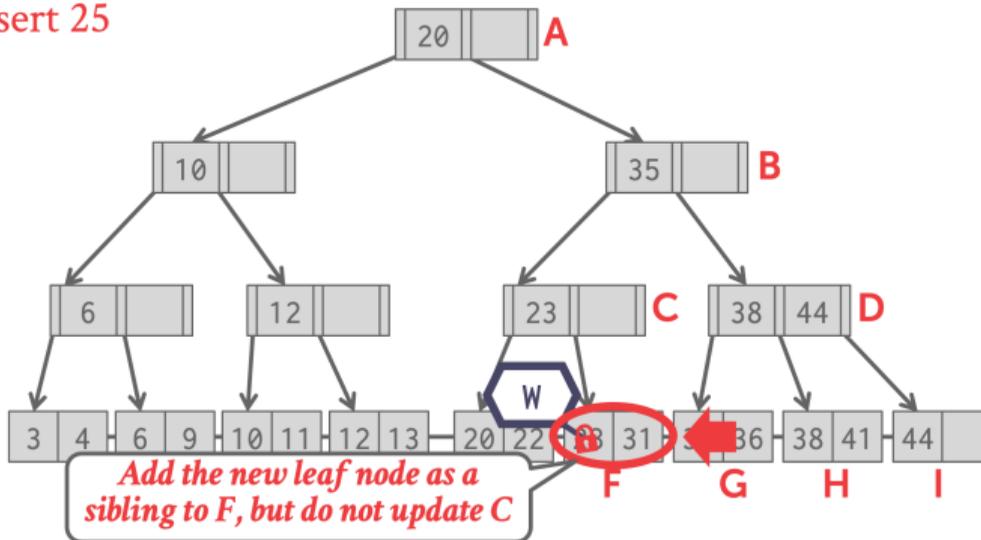


# Blink-Tree Example



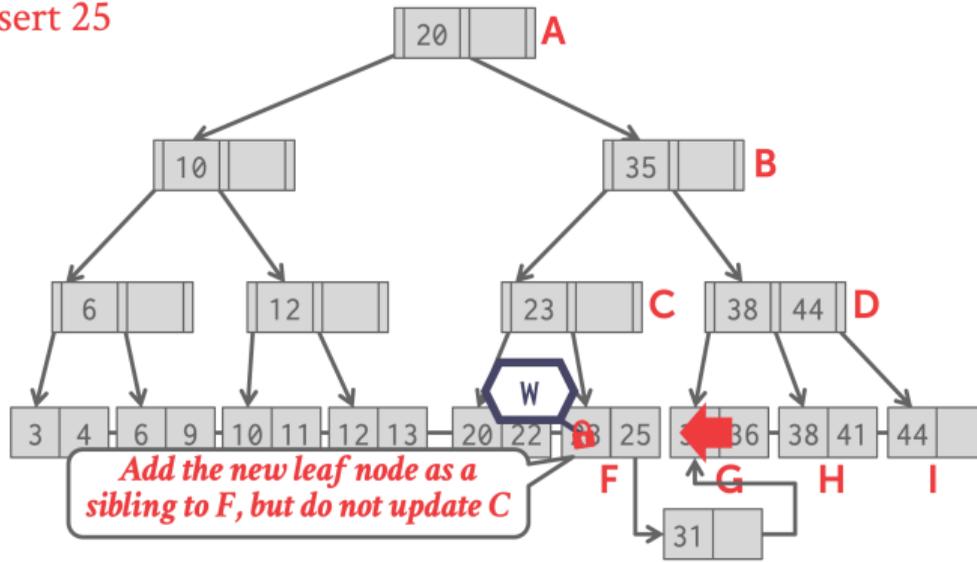
# B<sup>link</sup>-Tree Example

T<sub>1</sub>: Insert 25

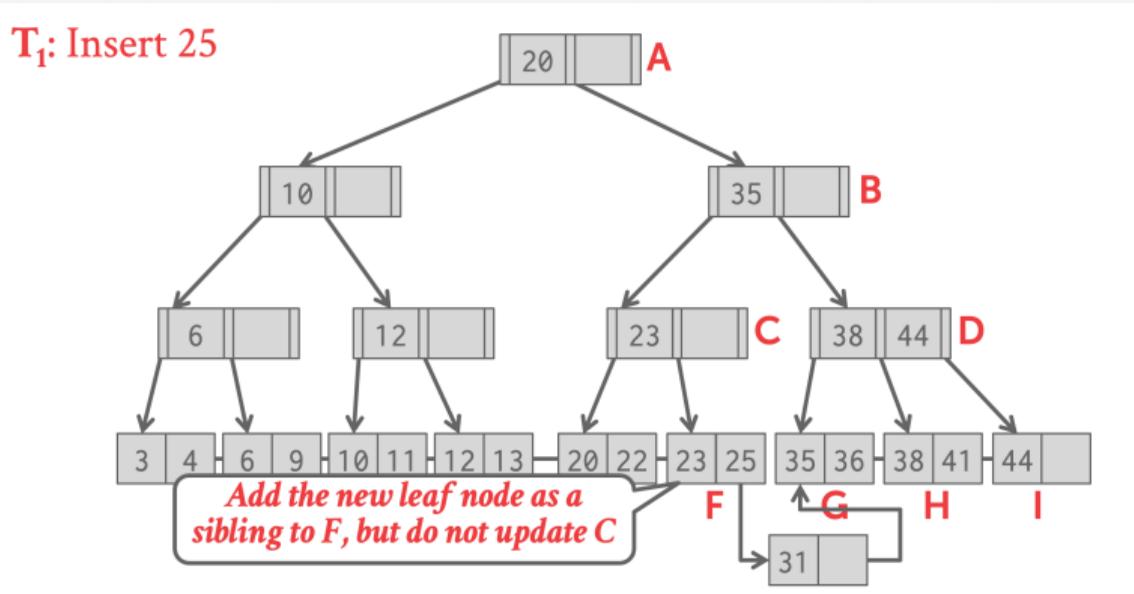


# B<sup>link</sup>-Tree Example

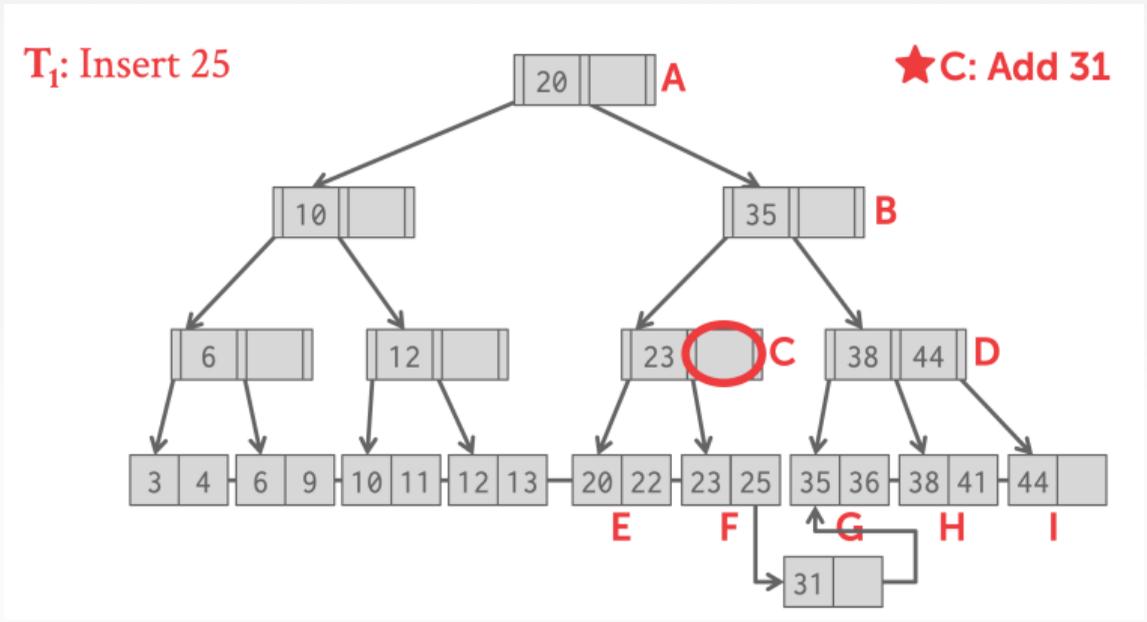
T<sub>1</sub>: Insert 25



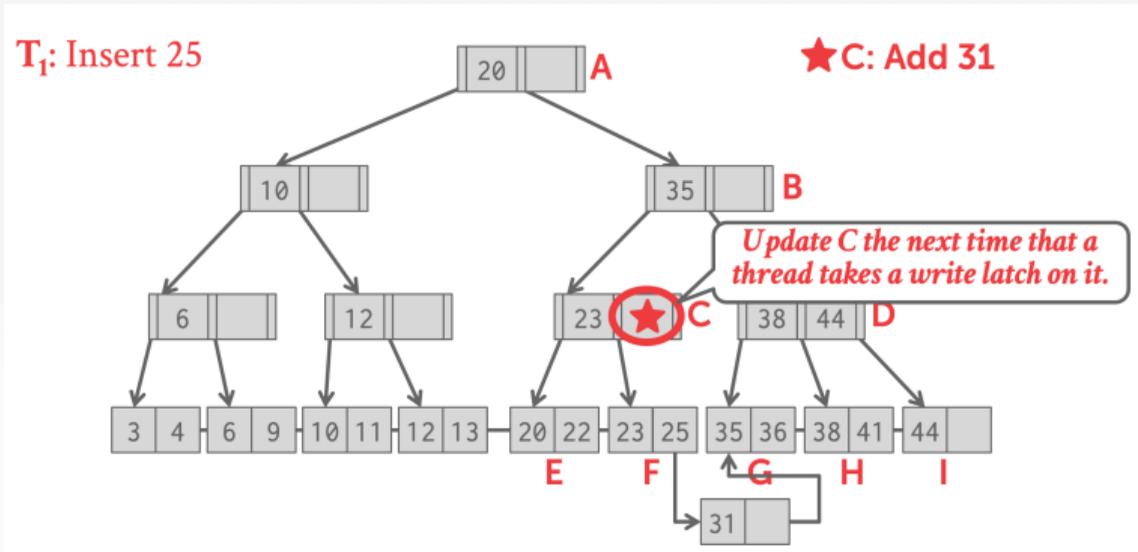
# Blink-Tree Example



# Blink-Tree Example



# Blink-Tree Example

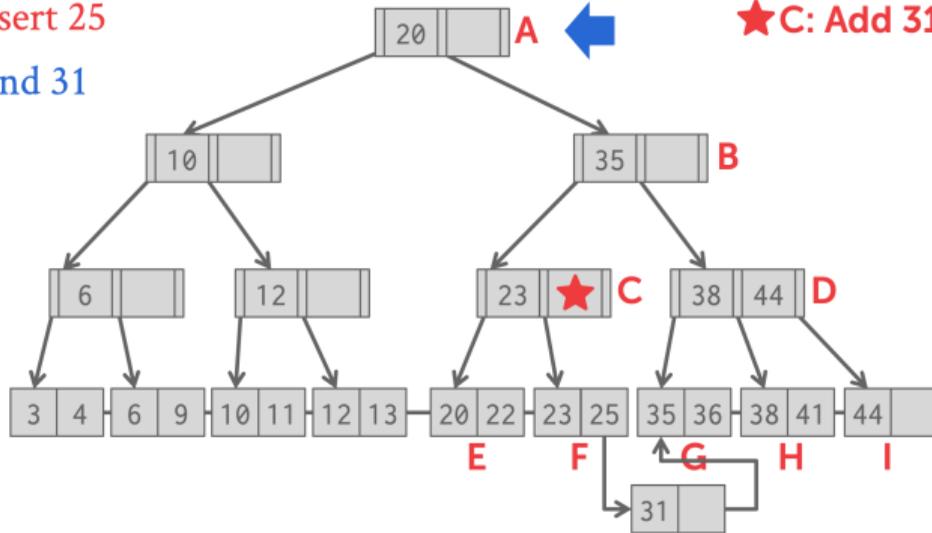


# Blink-Tree Example

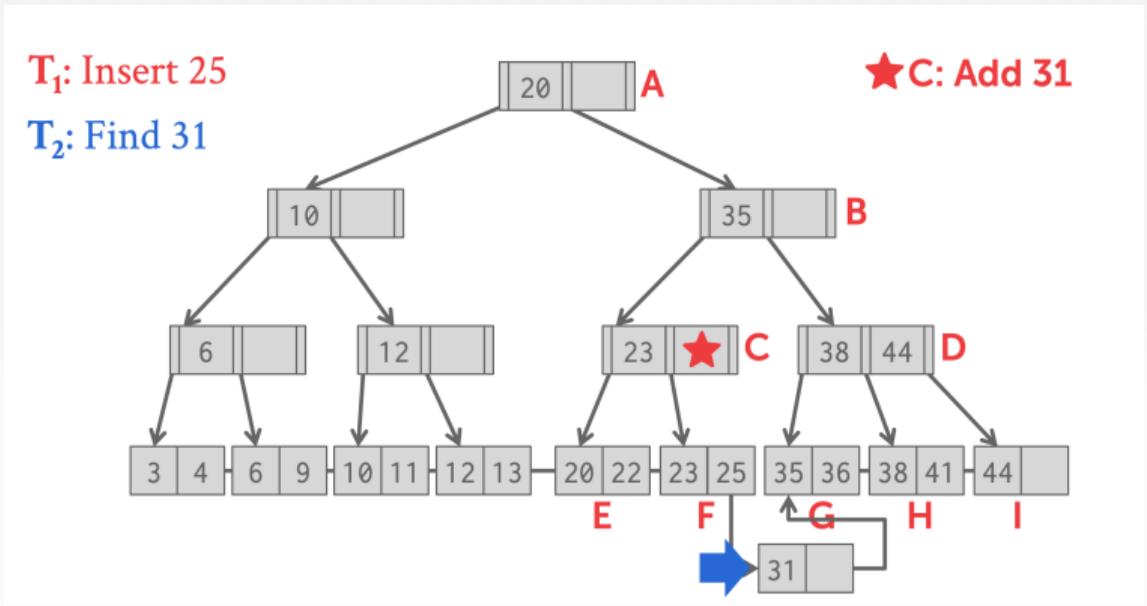
T<sub>1</sub>: Insert 25

T<sub>2</sub>: Find 31

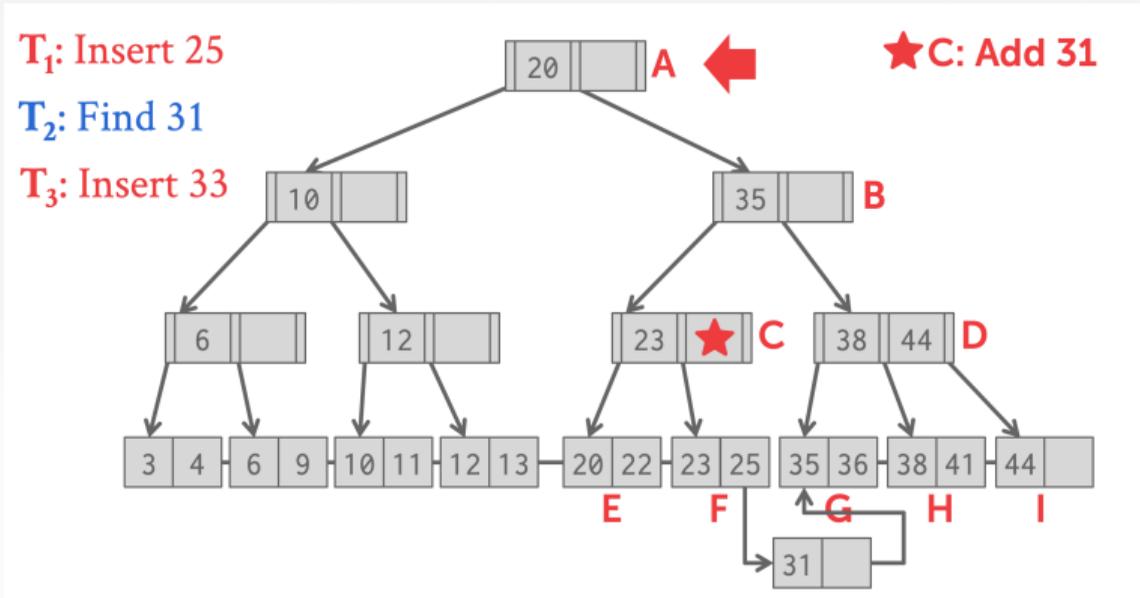
★ C: Add 31



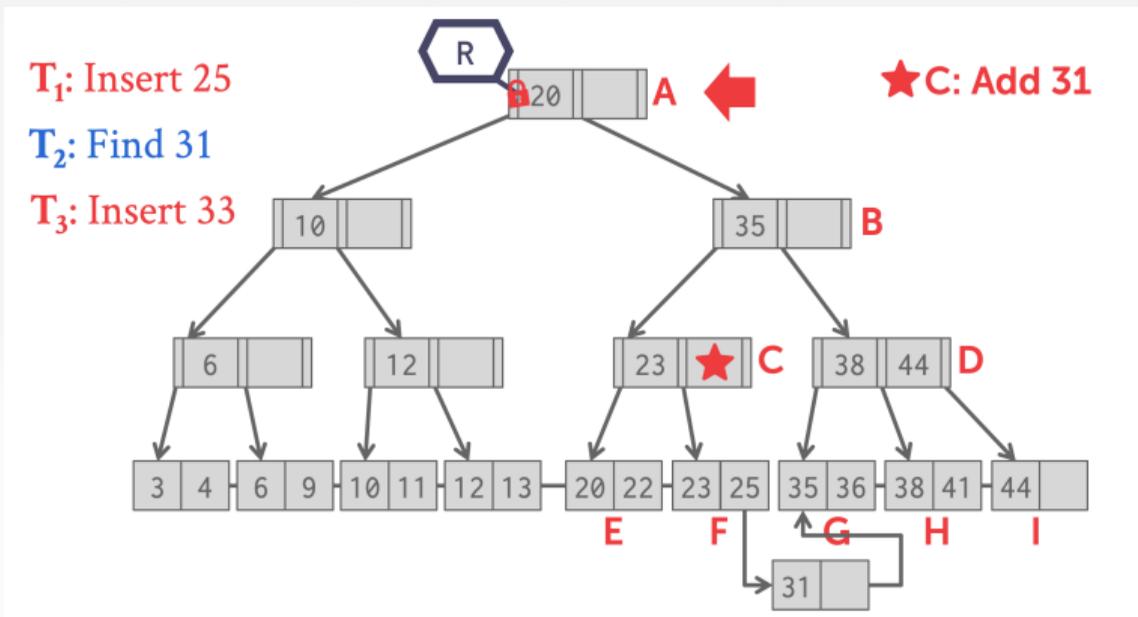
# B<sup>link</sup>-Tree Example



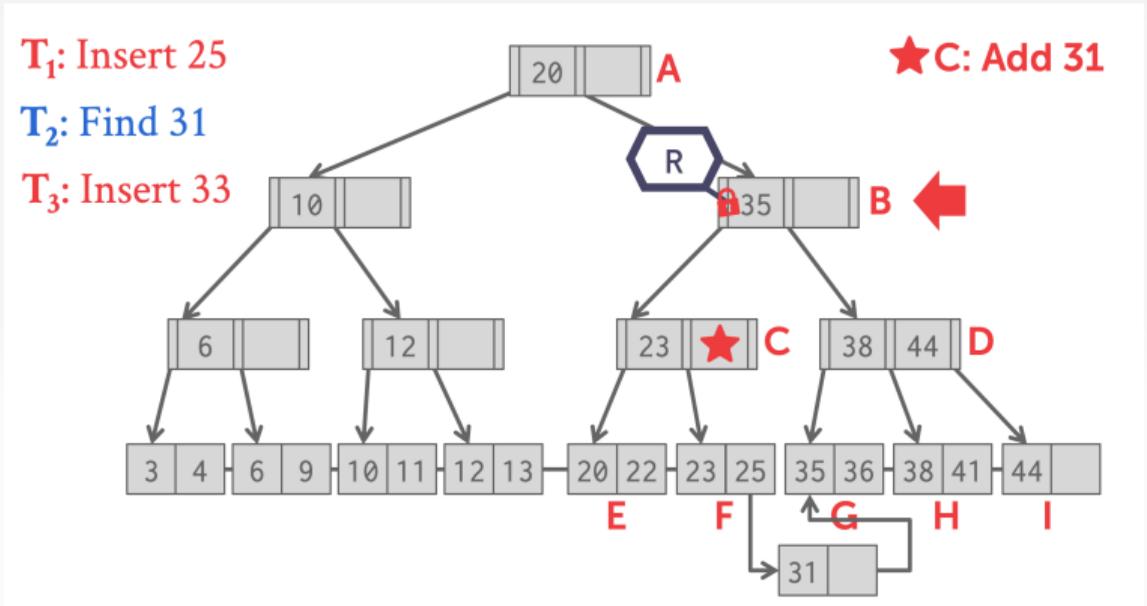
# Blink-Tree Example



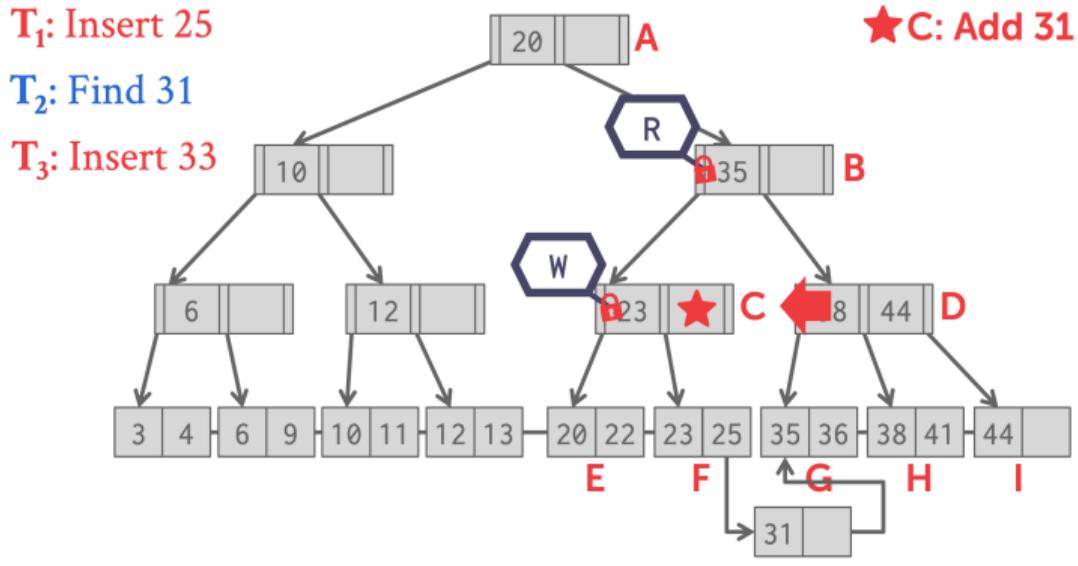
# Blink-Tree Example



# Blink-Tree Example



# Blink-Tree Example



# Blink-Tree Example

