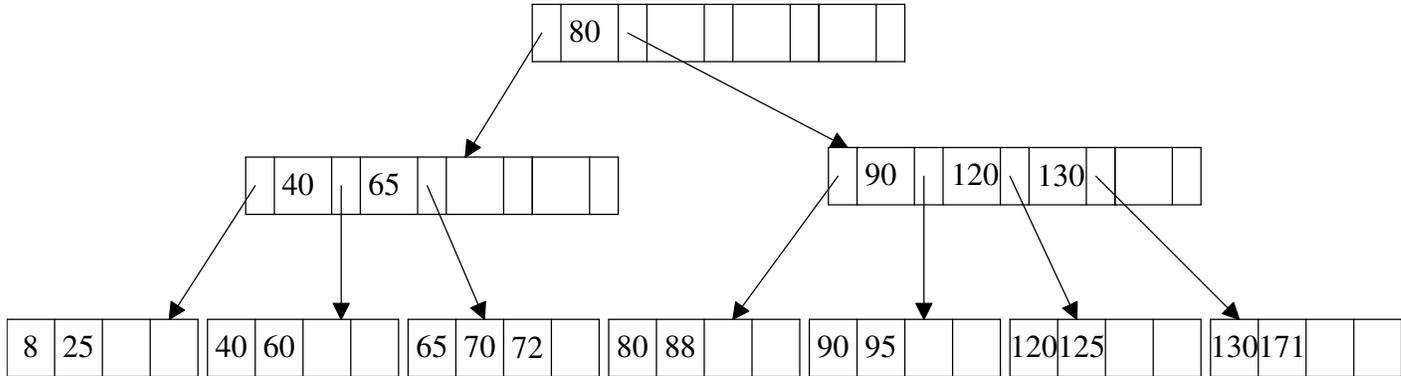


**Part A:** A B+ Tree is a multi-way tree (typically in the order of 100s children per node) used primarily as a file-index structure to allow fast search (as well as insertions and deletions) for a target key on disk. Two types of *pages* (B+ tree "nodes") exist:

- 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?

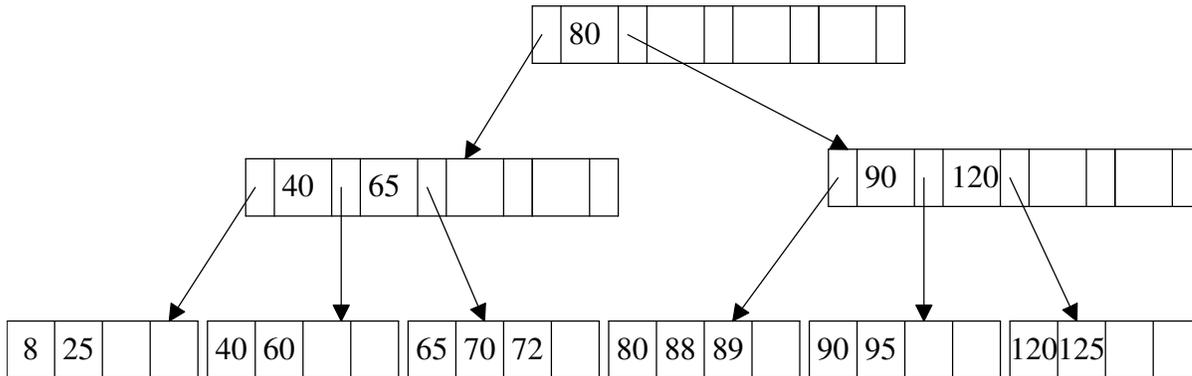
Situation		insertion Algorithm
Data Page Full?	Parent Index Page Full?	
No	No	Place record in sorted position in the appropriate data page.
Yes	No	<ol style="list-style-type: none"> <li>1. Split data page with records &lt; middle key going in left data page and records ≥ middle key going in right data 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 &lt; middle key going in left data page and records ≥ middle key going in right data page.</li> <li>2. Adding middle key to parent index page causes it to split with keys &lt; middle key going into the left index page, keys &gt; 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>

c) For a B+ tree with a branch factor 201, what would be the worst case height of the tree if the number of keys was 1,000,000,000,000?

**Part B:** The deletion algorithm for a B+ tree is summarized by the below table.

Situation		deletion Algorithm
Data Page Below Fill Factor?	Parent Index Page Below Fill Factor?	
No	No	Delete record from the data page. Shifting records with larger keys to left to fill in the hole. If the deleted key appears in the index page, use the next key to replace it.
Yes	No	1. Combine data page and its sibling. Change the index page to reflect the change.
Yes	Yes	1. Combine data page and its sibling. 2. Adjusting the index page to reflect the change causes it to drop below the fill factor, so combine the index page with its sibling. 3. Continue combining the next higher level index pages until you reach an index page with the correct fill factor or you reach the root index page.

Consider an B+ tree example with  $b = 5$  and 50% fill factor. Delete 89, 65, and 88. What is the resulting B+ tree?



If you have time in lab, finish old labs or start homework #6.