17.6 Summary

- BSTs have best-case height $\Theta(\log N)$, and worst-case height $\Theta(N)$.
- Big O is *not* the same as worst-case!
- B-Trees are a modification of the BST that maintain $\Theta(\log N)$ runtime for `add` and `contains` in the worst case. They maintain perfect balance during insertion.
- A B-Tree has a limit L on the number of values a node can hold, instead of having one item per node like a BST.
- Upon `add` in a B-Tree, we simply append the value to an existing leaf node in the correct location instead of creating a new leaf node. If the node is too full, it splits and pushes a value up.

Previous
17.5 B-Tree Performance

Next
17.7 Exercises

Last updated 1 year ago