

B-Trees

CS 225 - Fall 2025

Mattox Beckman / Based on slides from Brad Solomon

Objectives

- Explain how engineering reality can affect trees
- Explain how B Tree find and insert work

Announcements

- There were a LOT of cases for MP Lists
- Informal Early Feedback in the Discord

Pros:

Pros:

$\mathcal{O}(\log n)$ insert

Optimal 1D range queries

Pros:

$\mathcal{O}(\log n)$ insert

Optimal 1D range queries

Cons:

Pros:

$\mathcal{O}(\log n)$ insert

Optimal 1D range queries

Cons:


$\mathcal{O}(\log n)$ is actually slow

Needs a **LOT** of memory

Big-O assumes all operations take the same time

This is not always true!

Can you think of examples?



Can we always fit our data in main memory?

Where else can we keep our data?

Does this match our assumption that lookups are $\mathcal{O}(n)$?

Can we always fit our data in main memory?

- No! Just check your cell phone camera....

Where else can we keep our data?

Does this match our assumption that lookups are $\mathcal{O}(n)$?

Can we always fit our data in main memory?

- No! Just check your cell phone camera....

Where else can we keep our data?

- Hard Drives, the Cloud, Mass Storage....

Does this match our assumption that lookups are $\mathcal{O}(n)$?

Can we always fit our data in main memory?

- No! Just check your cell phone camera....

Where else can we keep our data?

- Hard Drives, the Cloud, Mass Storage....

Does this match our assumption that lookups are $\mathcal{O}(n)$?

- Nope. See <https://gist.github.com/hellerbarde/2843375>



Keep the number of seeks low

When possible store data locally

Make sure data is relevant

Keep the number of seeks low

- Make a tree that is wide and short

When possible store data locally

Make sure data is relevant

Keep the number of seeks low

- Make a tree that is wide and short

When possible store data locally

- Store more than one key per node

Make sure data is relevant

Keep the number of seeks low

- Make a tree that is wide and short

When possible store data locally

- Store more than one key per node


Make sure data is relevant

- Make sure the tree is ordered

A B Tree of order m is an m -ary tree

Nodes are **ordered**, have up to $m - 1$ keys and $keys + 1$ children

All leaves are on the same level



- Constructor
- Insert
- Find
- Delete

- Constructor
- Insert
- Find
- Delete

B Tree Node of Order m

```
1 struct BTreeNode {  
2     std::vector<DataPair> elements;  
3     std::vector<BTreeNode*> children;  
4 }
```

elements

-3	5	20	44
----	---	----	----

children

•	•	•	•	•
---	---	---	---	---

Base Case

- root empty \Rightarrow return
- leaf \Rightarrow do array find
-

Recursive Case


- Array find for first \geq match.
- Recurse to appropriate child

Base Case

- root empty \Rightarrow return
- leaf \Rightarrow do array find
-

Recursive Case

- Array find for first \geq match.
- Recurse to appropriate child




Base Case

- root empty \Rightarrow return
- leaf \Rightarrow do array find
-

Recursive Case

- Array find for first \geq match.
- Recurse to appropriate child

find(6)




Base Case

- root empty \Rightarrow return
- leaf \Rightarrow do array find
-

Recursive Case

- Array find for first \geq match.
- Recurse to appropriate child

find(6)




Base Case

- root empty \Rightarrow return
- leaf \Rightarrow do array find
-

Recursive Case

- Array find for first \geq match.
- Recurse to appropriate child

find(16)




Base Case

- root empty \Rightarrow return
- leaf \Rightarrow do array find
-

Recursive Case

- Array find for first \geq match.
- Recurse to appropriate child

find(16)




Base Case

- root empty \Rightarrow return
- leaf \Rightarrow do array find
-

Recursive Case

- Array find for first \geq match.
- Recurse to appropriate child

find(45)




Base Case

- root empty \Rightarrow return
- leaf \Rightarrow do array find
-

Recursive Case

- Array find for first \geq match.
- Recurse to appropriate child

find(45)



Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
-	-	-	-	-

Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
5	-	-	-	-

Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
3	5	-	-	-

Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
3	5	8	-	-

Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
2	3	5	8	-

Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
2	3	4	5	8

Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
2	3	4	5	8

What is the time complexity of sorted array insertion?

Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
2	3	4	5	8

What is the time complexity of sorted array insertion?

It's $\mathcal{O}(n)$. Is this a problem?

Node insert is simply ordered array insert.

Insert 5, 3, 8, 2, 4

0	1	2	3	4
2	3	4	5	8

What is the time complexity of sorted array insertion?

It's $\mathcal{O}(n)$. Is this a problem?

No. m is constant, and memory operations are **fast**.

What do we do if the node gets full?

Split the node and promote the median.


Let's suppose $m = 5$.

0	1	2	3	4
2	3	4	5	8

What do we do if the node gets full?

Split the node and promote the median.


Let's suppose $m = 5$.



Splitting can happen recursively.

Let $m = 3$. Insert 25.

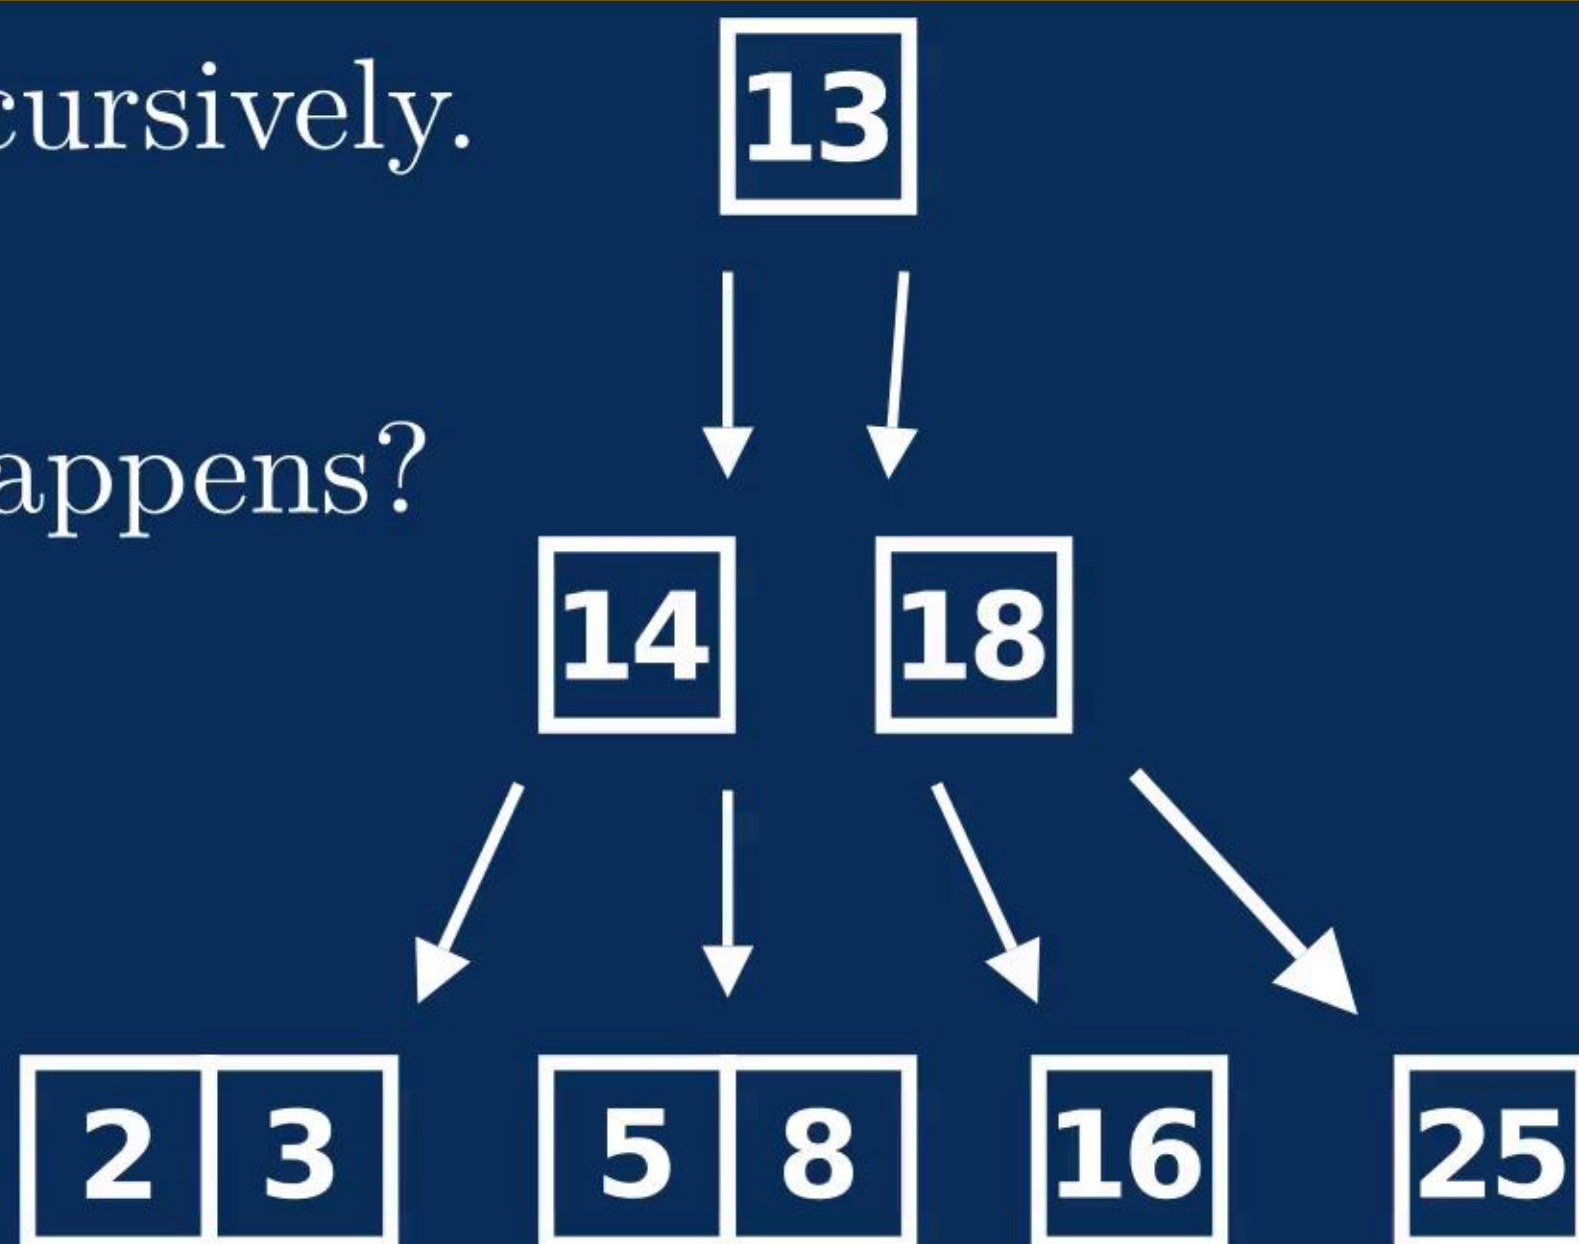
Can you predict what happens?



Splitting can happen recursively.

Let $m = 3$. Insert 25.

Can you predict what happens?



Next time: B Tree delete and analysis