

Data Structures

BTree

CS 225


October 4, 2023

Brad Solomon & G Carl Evans



UNIVERSITY OF
ILLINOIS
URBANA - CHAMPAIGN

Department of Computer Science



Exam 2 Summary → Very well!

Average: 87%

Median: 93%

30% of class got 100%

Most missed questions on exam 2 were concepts from exam 1

↳ Exams are cumulative!

↳ If you aren't performing as well as you would like...

Go to office hours (at off peak times)!

MP Mosaics Extra Credit closes today!

↳ Extended for this MP only

MP Art

```
=====  
All tests passed (409479 assertions in 35 test cases)  
root@cf914a57bfac:/workspaces/cs225git/mp_stickers/build#
```

INGBONES TUMBLR



BUT WHY **are** ALL THE **ERRORS** GONE?

MP Art

Submitted answer 1
Submitted at 2023-08-30 21:35:35 (CDT) **invalid, not gradable** show

- × [0/5] A StickerSheet with no stickers works
- × [0/5] StickerSheet does not duplicate stickers in memory
- × [0/5] StickerSheet behaves correctly when a sticker is updated

Save & Grade Unlimited attempts Save only

Grading possible in 238 min Can only be graded once every 360 minutes

- × [0/5] StickerSheet::setStickerAtLayer() fails for an invalid layer
- × [0/1] StickerSheet::changeMaxStickers() can increase the number

✓ [10/10] myImage.png exists and contains stickers

```
root@5d9182915a92:/workspaces/CS_225/cs225git/mp_stickers/build# sudo ./main  
Segmentation fault
```

MP Art



B/c mosaics has cool art too!



Learning Objectives

↗ Another tree

Discuss alternatives to AVL Trees (and BSTs)

Implement the BTree



Considering hardware limitations

Can we always fit our data in main memory?

↳ No!

Assume it

↳ Yes!

$X \rightarrow Y$

$O(1)$

Where else can we keep our data?

Hard drive

External drive

AWS / cloud / download :-

Does this match our assumption that all memory lookups are $O(1)$?

No!

BTree Design Motivations

Node
child is slow

(key, value)
↑
write


When large seek times become an issue, we address this by:

1. "Pack a node with more data" — **store many keys in each node**


2. we want to have a *n!b* tree

↳ More children for every node

↳ My tree height will be smaller




3. Make sure my tree is ordered/sorted! (Make data relevant)



only certain trees can do this (easily)

Store my tree as an array in local memory



BTree

$O(\log_2 n)$

of order 2 \rightarrow
2
 $\rightarrow \dots$



A BTree (of order m) is a m -ary tree

Nodes contain up to $m-1$ keys and have $|\text{keys}|+1$ children

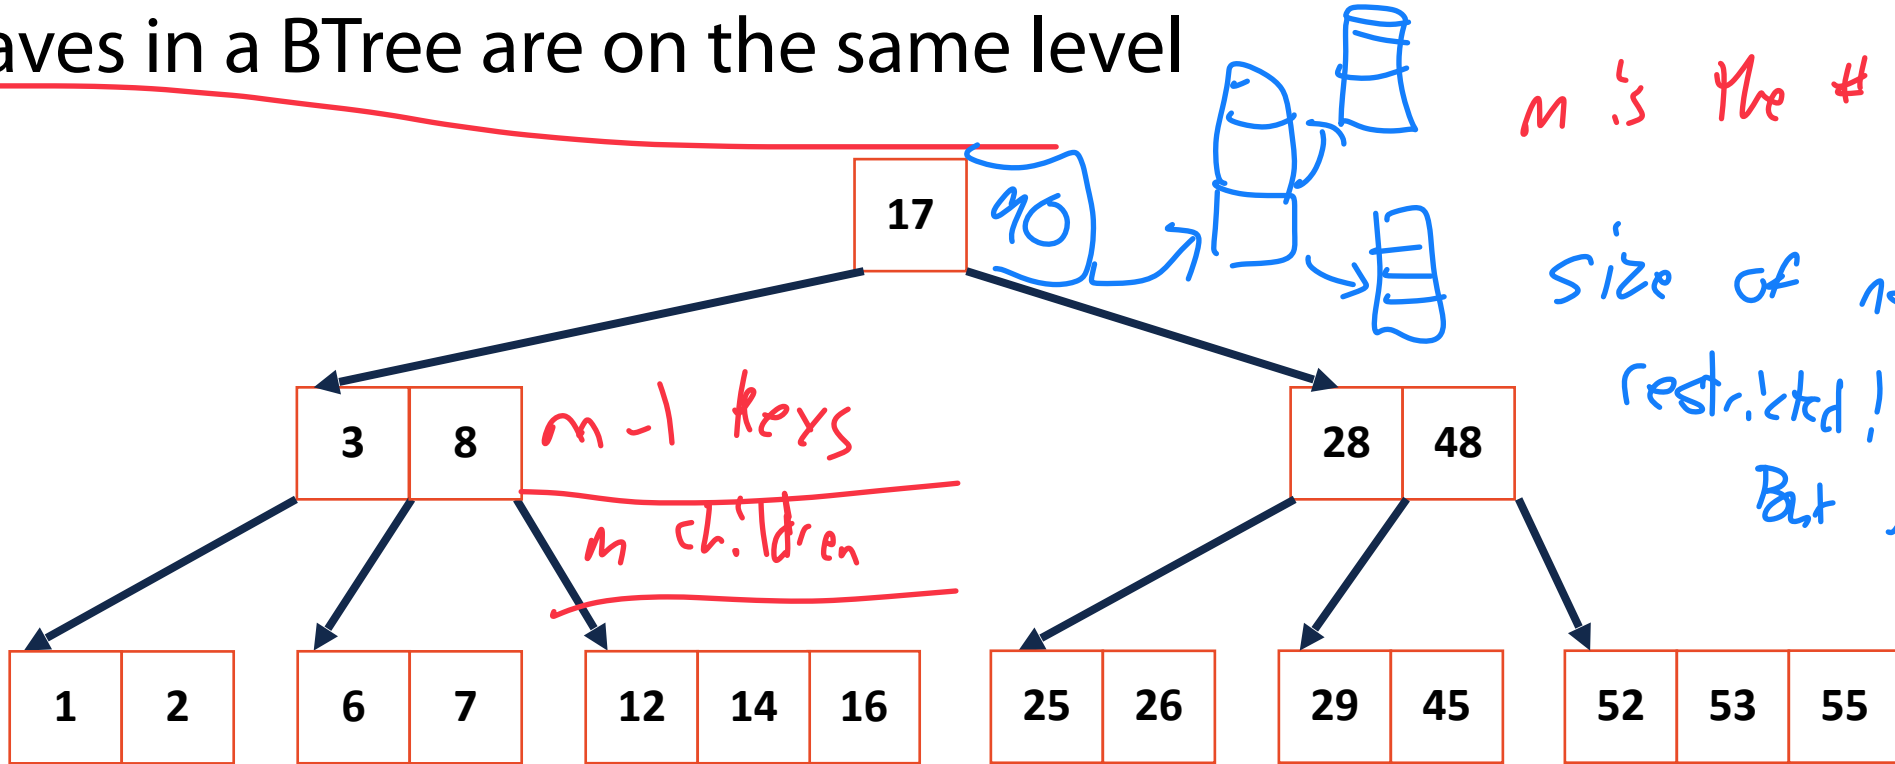
All leaves in a BTree are on the same level

m is the # of children

Size of nodes can be restricted!

But not by level.

why? \rightarrow



$m-1$ keys

m children

$m-1$ keys


$m \geq 4$

BTree Node (of order m)

$M \geq 5$



Smallest not keys largest



↳ an array of values

↳ an array of children



class BTree Node {

vector (key, value pair) objects;
↳ Data Pair


↳ vector (BTree Node*) children;

↓
Tradeoff!
↳ faster if precall everything
↳ cost is size!


what values



m different "regions"



↳ objects 1 per



↳ objects m per shard

BTree Node (of order m)

Motivation is memory

1 byte = 8 bits
4 bytes = 32 bits

What value of m should we be using?

↳ we want to determine this based on need or seek architecture slower

Disk Block

↳

RAM



↳ Let m equal at most one of these blocks

↳ 512 bytes or bits

↳ 1096 bytes

↳ int is 32 bits!

$$m = \left\{ \begin{array}{l} \text{(size of chunk)} \\ \text{Size of object} \end{array} \right.$$

sending info by cloud

TCP Network Packet $\rightarrow \sim 1500$ bytes


BTree Insertion

node (sorted)

All keys within a BTree are ordered

$$M = 5$$

↳ 4 keys
↳ 5 children



Insert (10)

Insert (5)

(5)





Array insert is slow!

↳ in a sorted array

BTree Insertion

All keys within a BTree are ordered



↳ too big!



Insert (10)

↳ 2 op

Insert (5)

↳ 2 op

Insert (7)

↳ 3 op

Insert (9)

↳ 4 op

Insert (2)


...

↳ We will split array

At most $M-1$
Allocate M

M = 5

Allocate M space for each node*




BTree Insertion

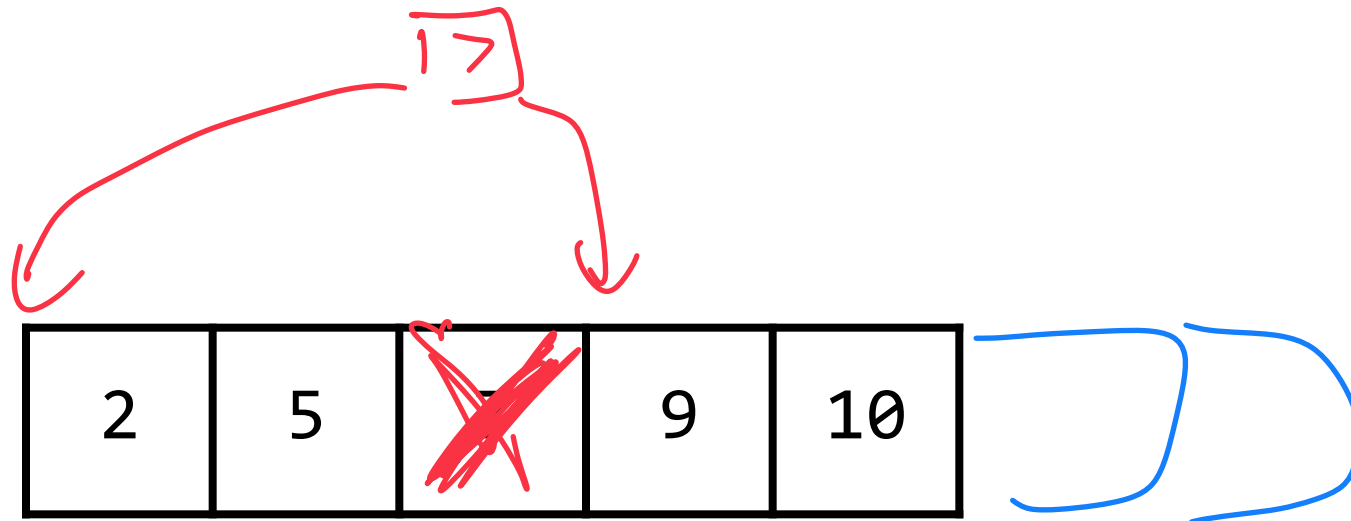

M = 5

When a BTree node reaches m keys:

- 1) Find median value
- 2) Split by "throwing up" the median
↳ we make a new parent



key 6
vector [vals] ~~No!~~




BTree Insertion

M = 5


When a BTree node reaches **m** keys, **split and make a new parent**.

1) Find median


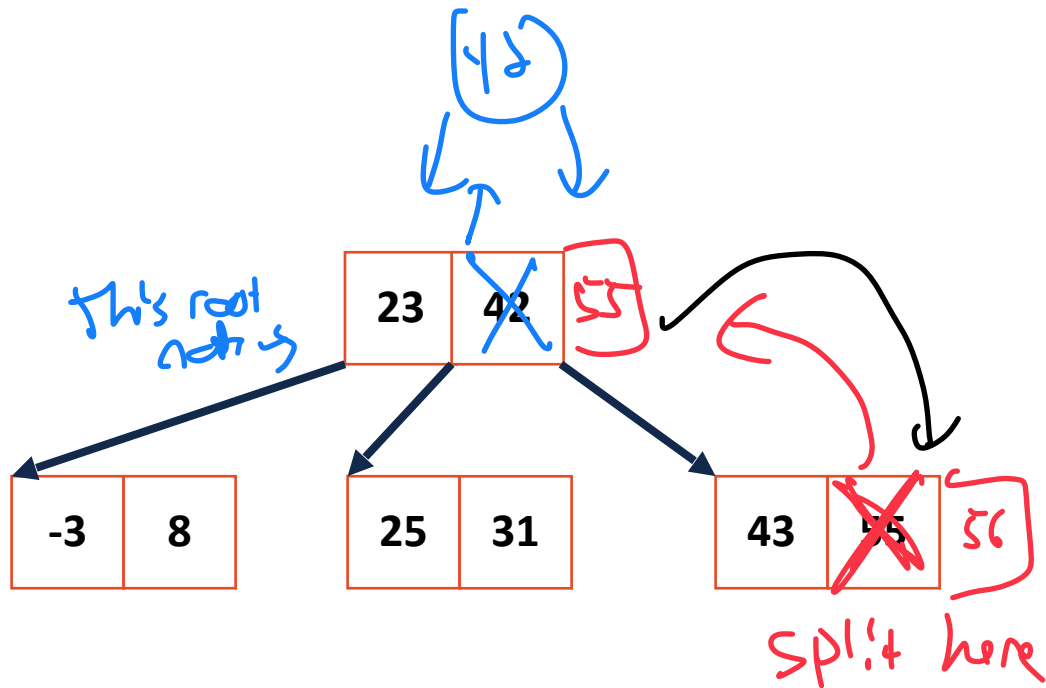


BTree Recursive Insert

Insert (56), M = 3



Insert always starts at a leaf but can propagate up repeatedly.



BTree Visualization/Tool




<https://www.cs.usfca.edu/~galles/visualization/BTree.html>

↑
play around w/ this visualizer!

BTree Size Restrictions

$m-1$ on keys (of order m)

By definition we have max, but do we have min? Are these trees valid?





BTree Properties

A **BTree** of order **m** is an m-ary tree and by definition:

- All keys within a node are ordered
- All nodes contain no more than **m-1** keys.
- All internal nodes have exactly **one more child than keys**


Root nodes can be a leaf or have _____ children.

All non-root, internal nodes have _____ children.

All leaves in the tree are at the same level.

BTree

If I tell you this is a valid BTree, what is the value of m ?





BTree ADT

Constructor


Insert

Find

Delete


BTree Find

Find(12)




BTree Find

Find(7)




BTree Find

Find(70)



BTree *Exists*


```
1 bool Btree::_exists(BTreeNode & node, const K & key) {
2
3     unsigned i;
4     for ( i = 0; i < node.keycount_ && key > node.keys_[i]; i++) { }
5
6     if ( i < node.keycount_ && key == node.keys_[i] ) {
7         return true;
8     }
9
10    if ( node.isLeaf() ) {
11        return false;
12    } else {
13        BTreeNode nextChild = node._fetchChild(i);
14        return _exists(nextChild, key);
15    }
16 }
```



BTree Exists




```
1 bool Btree::_exists(BTreeNode & node, const K & key) {
2
3     unsigned i;
4     for ( i = 0; i < node.keycount_ && key > node.keys_[i]; i++) { }
5
6     if ( i < node.keycount_ && key == node.keys_[i] ) {
7         return true;
8     }
9
10    if ( node.isLeaf() ) {
11        return false;
12    } else {
13        BTreeNode nextChild = node._fetchChild(i);
14        return _exists(nextChild, key);
15    }
16 }
```



BTree Remove


BTree removal is complicated! **It won't be part of the lab.**

However lets consider how we would handle the following cases...




BTree Remove

Remove (8)




BTree Remove

Remove (2)




BTree Remove

Remove (15)




BTree Remove

Remove (42)



BTree Remove

Remove (5)



For next time: BTree Analysis

We've seen the ADT

What is the runtime for our BTree operations?