



Priority Queues (Heaps)



Motivation

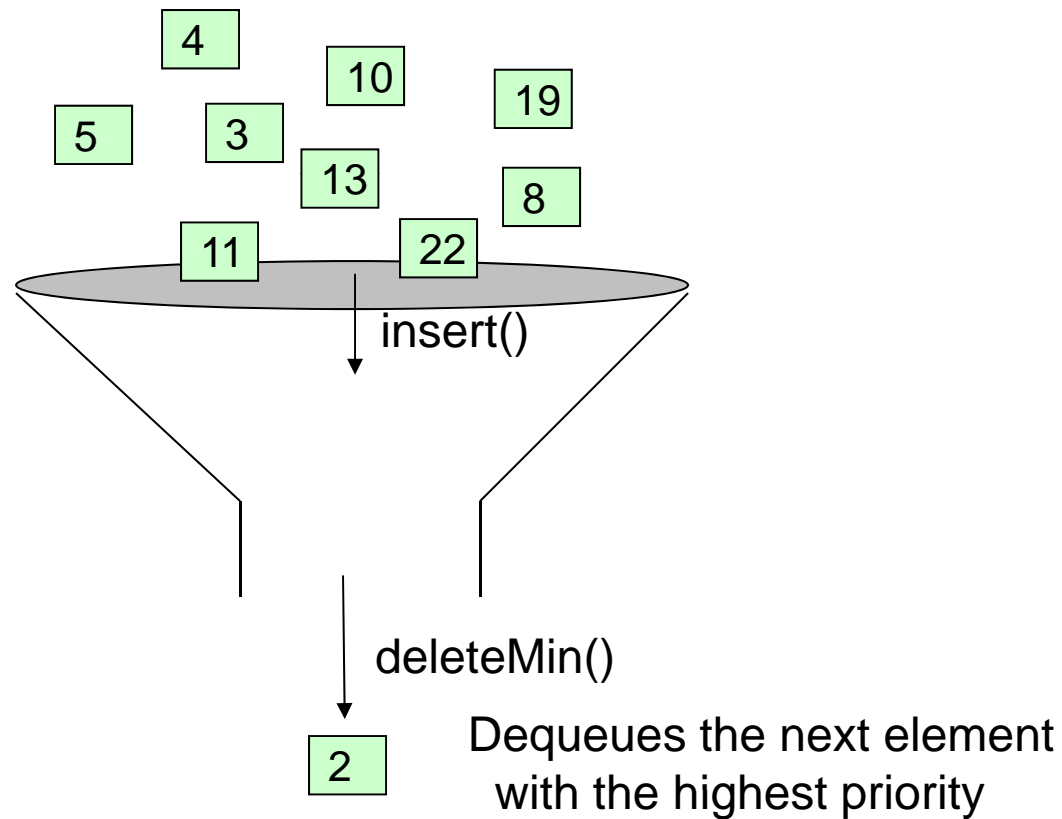
- Queues are a standard mechanism for ordering tasks on a first-come, first-served basis
- However, some tasks may be more important or timely than others (higher priority)
- Priority queues
 - Store tasks using a partial ordering based on priority
 - Ensure highest priority task at head of queue
- Heaps are the underlying data structure of priority queues



Priority Queues: Specification

- Main operations
 - **insert** (i.e., enqueue)
 - Dynamic insert
 - specification of a priority level (0-high, 1,2.. Low)
 - **deleteMin** (i.e., dequeue)
 - Finds the current minimum element (read: “highest priority”) in the queue, deletes it from the queue, and returns it
- Performance goal is for operations to be “fast”

Using priority queues




Can we build a data structure better suited to store and retrieve priorities?

Simple Implementations


- Unordered linked list

- $O(1)$ insert
- $O(n)$ deleteMin



- Ordered linked list

- $O(n)$ insert
- $O(1)$ deleteMin




- Ordered array


- $O(\lg n + n)$ insert
- $O(n)$ deleteMin



- Balanced BST

- $O(\log_2 n)$ insert and deleteMin





Binary Heap

A priority queue data structure



Binary Heap

- A binary heap is a binary tree with two properties
 - Structure property
 - Heap-order property




Structure Property

- A binary heap is a complete binary tree
 - Each level (except possibly the bottom most level) is completely filled
 - The bottom most level may be partially filled (from left to right)

- Height of a complete binary tree with N elements is $\lfloor \log_2 N \rfloor$

Binary Heap Example

N=10



Array representation:


	A	B	C	D	E	F	G	H	I	J			
0	1	2	3	4	5	6	7	8	9	10	11	12	13



Heap-order Property

- Heap-order property (for a “MinHeap”)
 - For every node X , $\text{key}(\text{parent}(X)) \leq \text{key}(X)$
 - Except root node, which has no parent
- Thus, minimum key always at root
 - Alternatively, for a “MaxHeap”, always keep the maximum key at the root
- Insert and deleteMin must maintain heap-order property


Heap Order Property



- Duplicates are allowed
- No order implied for elements which do not share ancestor-descendant relationship

Implementing Complete Binary Trees as Arrays

- Given element at position i in the array
 - Left child(i) = at position $2i$
 - Right child(i) = at position $2i + 1$
 - Parent(i) = at position $\lfloor i/2 \rfloor$



```

1  template <typename Comparable>
2  class BinaryHeap
3  {
4  public:
5      explicit BinaryHeap( int capacity = 100 );
6      explicit BinaryHeap( const vector<Comparable> & items );
7
8      bool isEmpty( ) const;
9      const Comparable & findMin( ) const;
10
11     void insert( const Comparable & x );
12     void deleteMin( );
13     void deleteMin( Comparable & minItem );
14     void makeEmpty( );
15
16 private:
17     int         currentSize; // Number of elements in heap
18     vector<Comparable> array; // The heap array
19
20     void buildHeap( );
21     void percolateDown( int hole );
22 };

```

Just finds the Min without deleting it

insert

deleteMin

Note: a general delete() function is not as important for heaps but could be implemented

Stores the heap as a vector

Fix heap after deleteMin



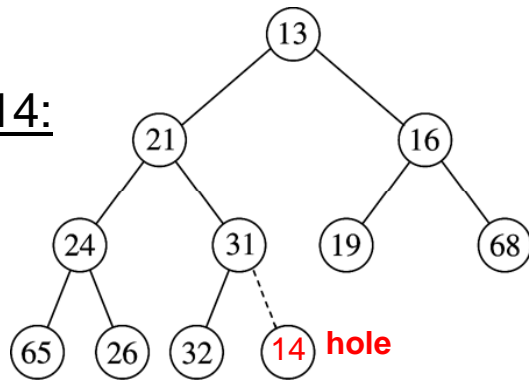
Heap Insert

- Insert new element into the heap at the next available slot (“hole”)
 - According to maintaining a complete binary tree
- Then, “percolate” the element up the heap while heap-order property not satisfied

Percolating Up


Heap Insert: Example

Insert 14:




Percolating Up

Heap Insert: Example



Heap Insert: Example


Insert 14:



Percolating Up

Heap Insert: Example

Insert 14:





Heap Insert: Implementation

```
// assume array implementation  
void insert( const Comparable &x) {  
?  
}
```



Heap Insert: Implementation

```
1    /**
2     * Insert item x, allowing duplicates.
3     */
4    void insert( const Comparable & x )
5    {
6        if( currentSize == array.size( ) - 1 )
7            array.resize( array.size( ) * 2 );
8
9        // Percolate up
10       int hole = ++currentSize;
11       for( ; hole > 1 && x < array[ hole / 2 ]; hole /= 2 )
12           array[ hole ] = array[ hole / 2 ];
13       array[ hole ] = x;
14   }
```

$O(\log N)$ time




Heap DeleteMin

- Minimum element is always at the root
- Heap decreases by one in size
- Move last element into hole at root
- *Percolate down* while heap-order property not satisfied

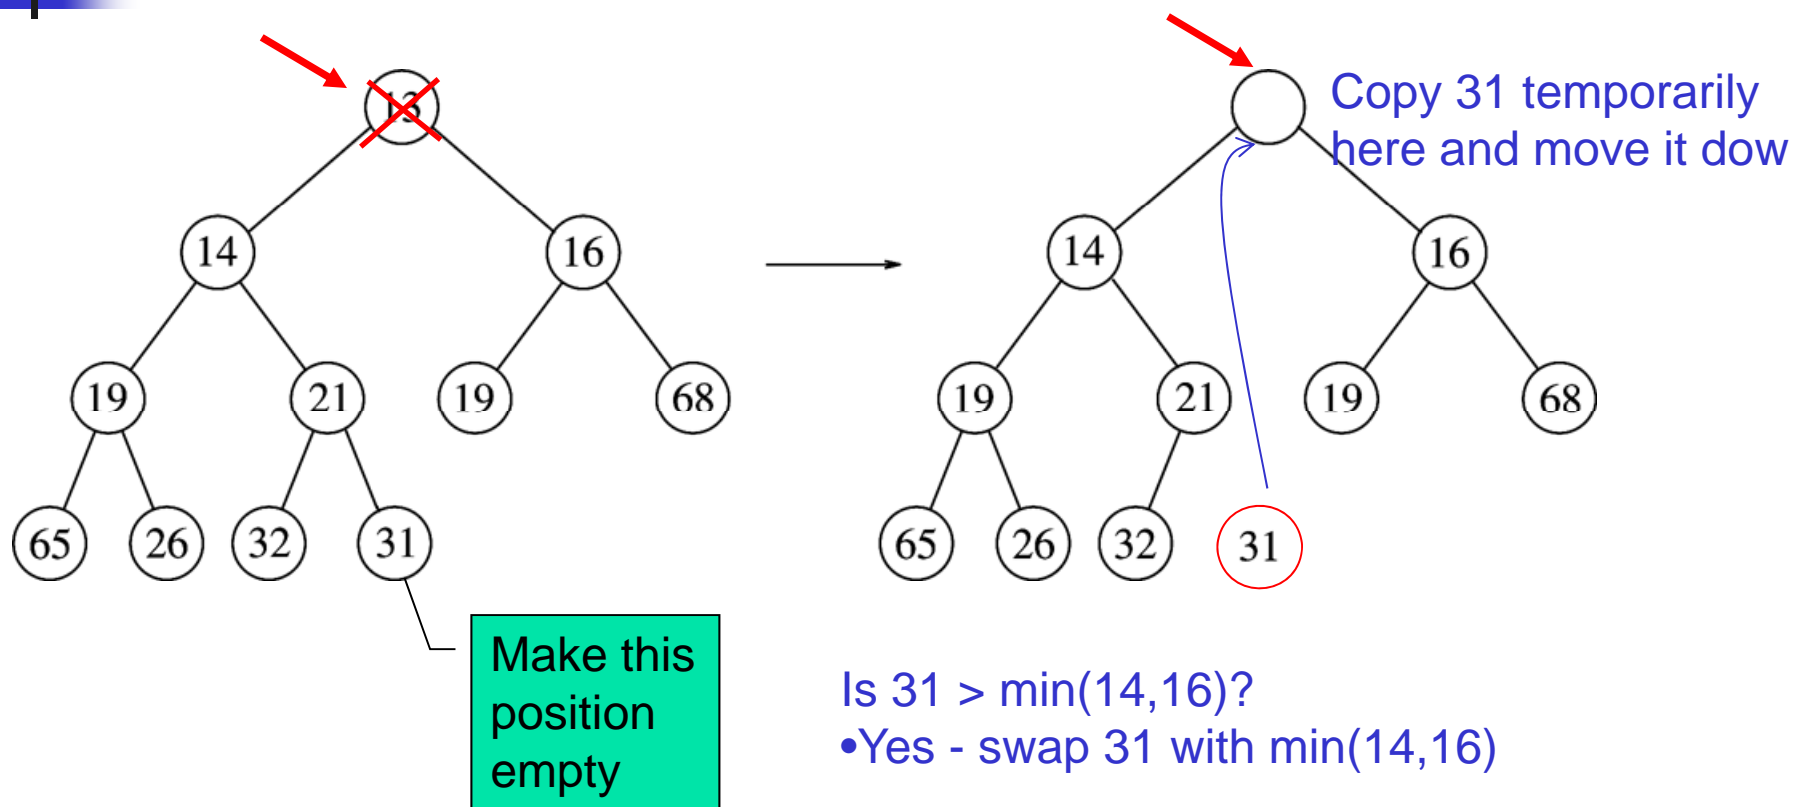
Percolating down...

Heap DeleteMin: Example




Percolating down...

Heap DeleteMin: Example



Percolating down...

Heap DeleteMin: Example




Is $31 > \min(19, 21)$?

- Yes - swap 31 with $\min(19, 21)$

Percolating down...

Heap DeleteMin: Example



Is $31 > \min(19, 21)$?

• Yes - swap 31 with $\min(19, 21)$

Is $31 > \min(65, 26)$?


• Yes - swap 31 with $\min(65, 26)$

Percolating down...

Cpt S 223. School of EECS, WSU

Percolating down...


Heap DeleteMin: Example



Percolating down...
Cpt S 223. School of EECS, WSU

Percolating down...

Heap DeleteMin: Example



- ✓ Heap order prop
- ✓ Structure prop

Heap DeleteMin: Implementation

```
1  /**
2   * Remove the minimum item.
3   * Throws UnderflowException if empty.
4   */
5  void deleteMin( )
6  {
7      if( isEmpty( ) )
8          throw UnderflowException( );
9
10     array[ 1 ] = array[ currentSize-- ];
11     percolateDown( 1 );
12 }
```

```
14  /**
15   * Remove the minimum item and place it in minItem.
16   * Throws UnderflowException if empty.
17   */
18  void deleteMin( Comparable & minItem )
19  {
20     if( isEmpty( ) )
21         throw UnderflowException( );
22
23     minItem = array[ 1 ];
24     array[ 1 ] = array[ currentSize-- ];
25     percolateDown( 1 );
26 }
```

$O(\log N)$ time

Heap DeleteMin: Implementation

```
28  /**
29   * Internal method to percolate down in the heap.
30   * hole is the index at which the percolate begins.
31   */
32  void percolateDown( int hole )
33  {
34      int child;
35      Comparable tmp = array[ hole ];
36
37      for( ; hole * 2 <= currentSize; hole = child )
38      {
39          child = hole * 2;
40          if( child != currentSize && array[ child + 1 ] < array[ child ] )
41              child++;
42          if( array[ child ] < tmp )
43              array[ hole ] = array[ child ];
44          else
45              break;
46      }
47      array[ hole ] = tmp;
48  }
```

Percolate
down

Left child

Right child

Pick child to
swap with



Other Heap Operations

- *decreaseKey(p, v)*
 - Lowers the current value of item p to new priority value v
 - Need to percolate up
 - E.g., promote a job
- *increaseKey(p, v)*
 - Increases the current value of item p to new priority value v
 - Need to percolate down
 - E.g., demote a job
- *remove(p)*
 - First, decreaseKey(p, $-\infty$)
 - Then, deleteMin
 - E.g., abort/cancel a job

Run-times for all three functions?

$O(\lg n)$



Improving Heap Insert Time

- What if all N elements are all available upfront?
- To build a heap with N elements:
 - Default method takes $O(N \lg N)$ time
 - We will now see a new method called `buildHeap()` that will take $O(N)$ time - i.e., optimal




Building a Heap

- Construct heap from initial set of N items
- Solution 1
 - Perform N inserts
 - $O(N \log_2 N)$ worst-case
- Solution 2 (use *buildHeap()*)
 - Randomly populate initial heap with structure property
 - Perform a percolate-down from each internal node ($H[\text{size}/2]$ to $H[1]$)
 - To take care of heap order property

BuildHeap Example

Input: { 150, 80, 40, 10, 70, 110, 30, 120, 140, 60, 50, 130, 100, 20, 90 }

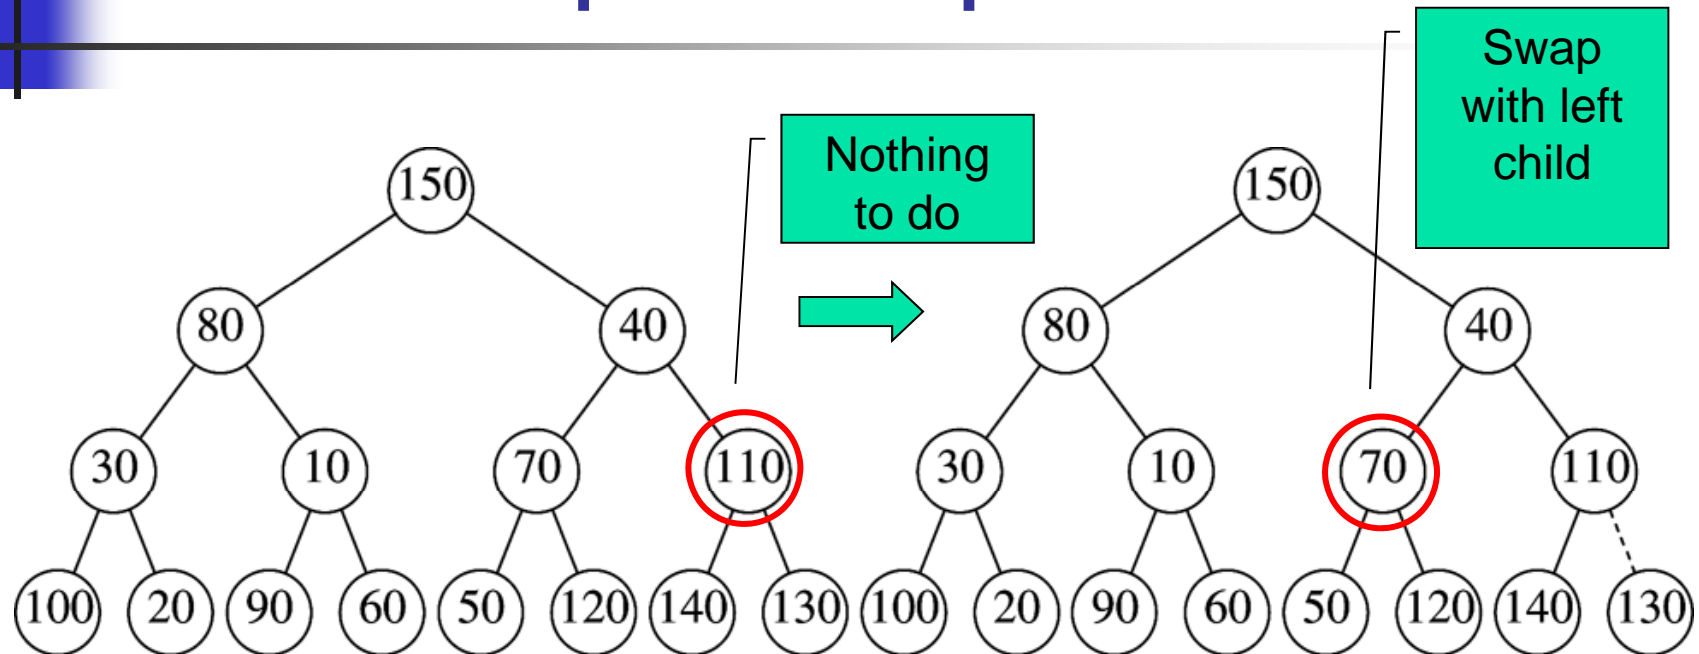


Leaves are all valid heaps (implicitly)

So, let us look at each internal node, from bottom to top, and fix if necessary

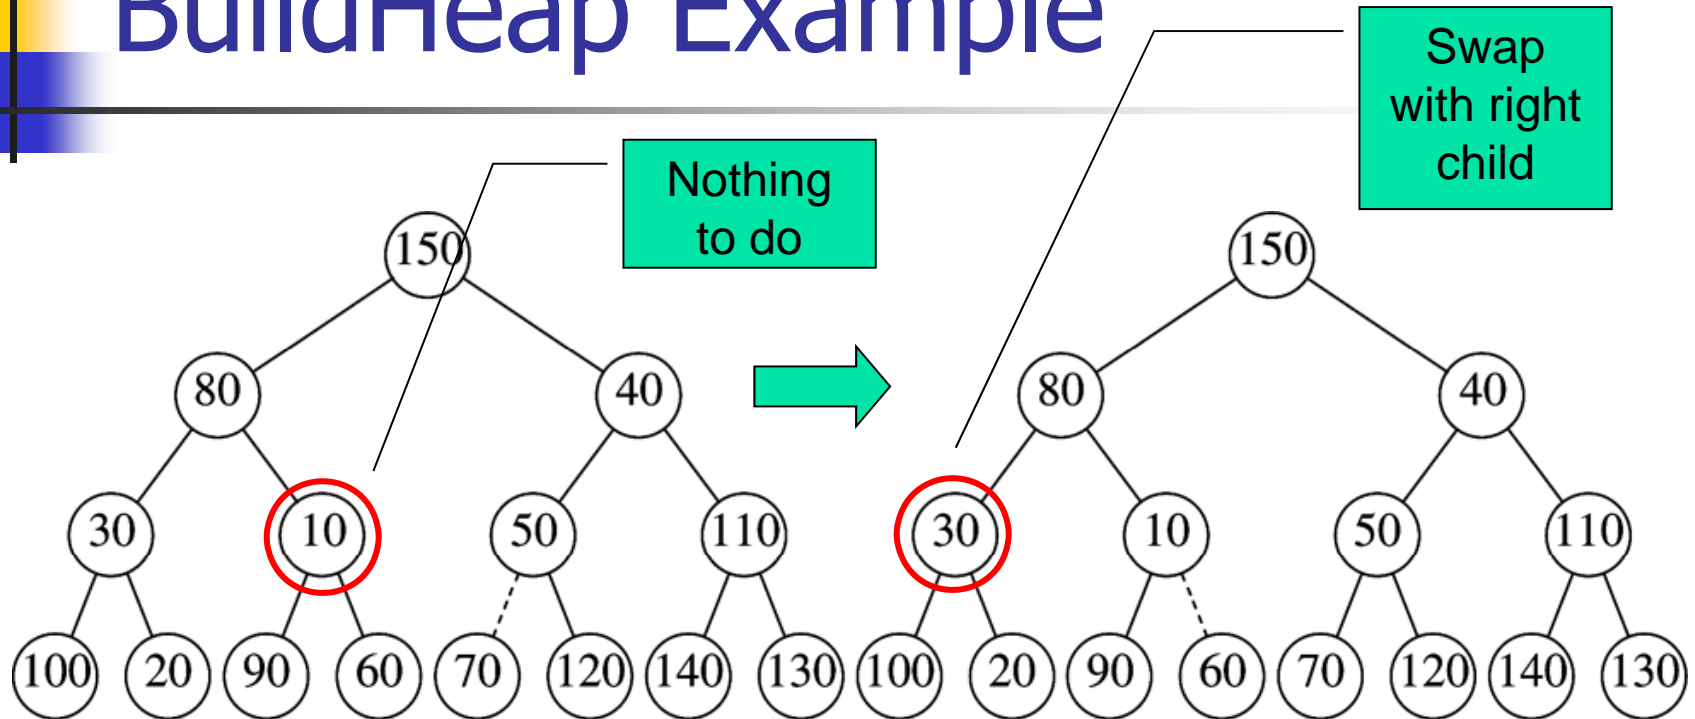
- Arbitrarily assign elements to heap nodes
- Structure property satisfied
- Heap order property violated
- Leaves are all valid heaps (implicit)

BuildHeap Example



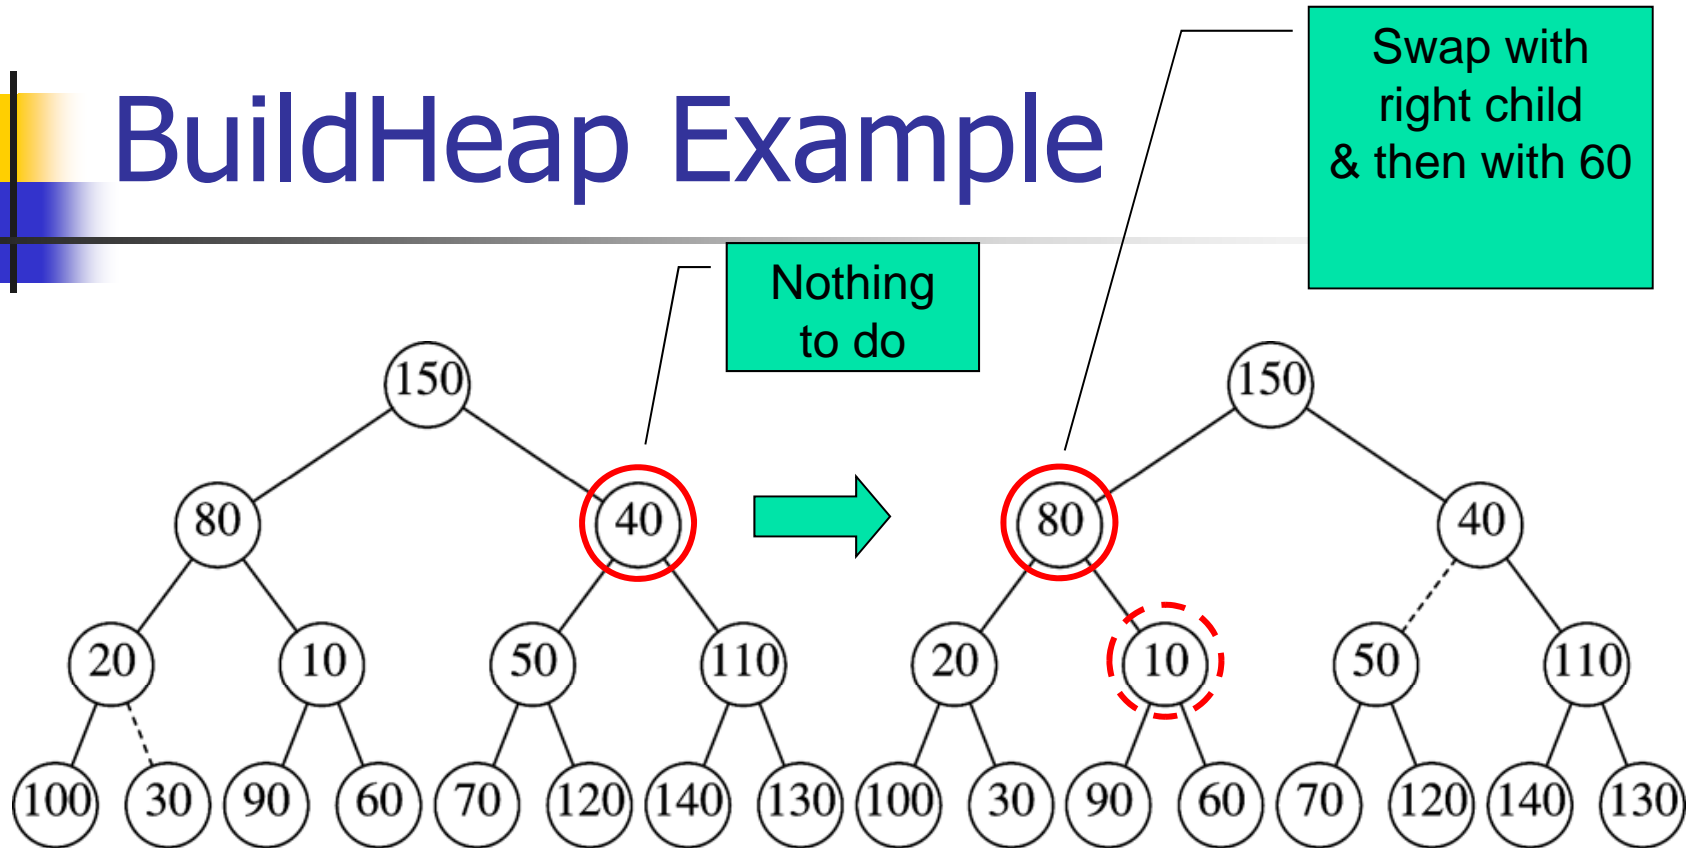
- Randomly initialized heap
- Structure property satisfied
- Heap order property violated
- Leaves are all valid heaps (simple)

BuildHeap Example




Dotted lines show path of percolating down

BuildHeap Example



Dotted lines show path of percolating down

BuildHeap Example



Final Heap

Dotted lines show path of percolating down

BuildHeap Implementation

```
1  explicit BinaryHeap( const vector<Comparable> & items )
2      : array( items.size( ) + 10 ), currentSize( items.size( ) )
3  {
4      for( int i = 0; i < items.size( ); i++ )
5          array[ i + 1 ] = items[ i ];
6      buildHeap( );
7  }
8
9  /**
10     * Establish heap order property from an arbitrary
11     * arrangement of items. Runs in linear time.
12     */
13  void buildHeap( )
14  {
15      for( int i = currentSize / 2; i > 0; i-- )
16          percolateDown( i );
17  }
```

Start with
lowest,
rightmost
internal node

BuildHeap() : Run-time Analysis

- Run-time = ?
 - **O(sum of the heights of all the internal nodes)**
because we may have to percolate all the way down to fix every internal node in the worst-case
- Theorem 6.1 **HOW?**
 - *For a perfect binary tree of height h , the sum of heights of all nodes is $2^{h+1} - 1 - (h + 1)$*
- Since $h = \lg N$, then sum of heights is $O(N)$
- Will be slightly better in practice


Implication: Each insertion costs $O(1)$ amortized time

Build heap analysis

Thm: Runtime $\equiv O(\text{sum of heights of all int nodes})$
 $\equiv O(\text{sum of heights of all nodes})$ incl. leaves
 $\equiv O(N)$

Proof:

~~Runtime~~ Sum of heights of all nodes = sum of heights of nodes at each level



$\sum_{i=0}^{h-1} \text{heights of all nodes}$

$$= \sum_{i=0}^{h-1} (h-i) \times 2^i$$

$$= \sum_{i=0}^{h-1} h 2^i - \sum_{i=0}^{h-1} i 2^i$$

$$= h \sum_{i=0}^{h-1} 2^i - \sum_{i=0}^{h-1} i 2^i$$

$$= h(2^{h+1} - 1) - S$$

$$= h(2^{h+1}) - h - (h-1)2^{h+1} + 2$$

$$= h \cdot 2^{h+1} - h + 2$$

$$= 2^{h+1} - h + 2 \quad \checkmark$$

As $h = \lg N$

$$\Rightarrow \sum \text{heights} = 2^{\lg N + 1} - \lg N + 2$$

$$= 2N - \lg N + 2 = O(N)$$



Binary Heap Operations

Worst-case Analysis

- Height of heap is $\lfloor \log_2 N \rfloor$
- insert: $O(\lg N)$ for each insert
 - In practice, expect less
- buildHeap insert: $O(N)$ for N inserts
- deleteMin: $O(\lg N)$
- decreaseKey: $O(\lg N)$
- increaseKey: $O(\lg N)$
- remove: $O(\lg N)$



Applications

- Operating system scheduling
 - Process jobs by priority
- Graph algorithms
 - Find shortest path
- Event simulation
 - Instead of checking for events at each time click, look up next event to happen



An Application: The Selection Problem

- Given a list of n elements, find the k^{th} smallest element

- Algorithm 1:
 - Sort the list $\Rightarrow O(n \log n)$
 - Pick the k^{th} element $\Rightarrow O(1)$

- A better algorithm:
 - Use a binary heap (minheap)



Selection using a MinHeap

- **Input: n elements**

- **Algorithm:**

- | | | |
|----|---|-------------------|
| 1. | buildHeap(n) | ==> $O(n)$ |
| 2. | Perform k deleteMin() operations | ==> $O(k \log n)$ |
| 3. | Report the k^{th} deleteMin output | ==> $O(1)$ |

Total run-time = $O(n + k \log n)$

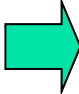

If $k = O(n/\log n)$ then the run-time becomes $O(n)$



Other Types of Heaps

- Binomial Heaps
- d-Heaps
 - Generalization of binary heaps (ie., 2-Heaps)
- Leftist Heaps
 - Supports merging of two heaps in $o(m+n)$ time (ie., sub-linear)
- Skew Heaps
 - $O(\log n)$ amortized run-time
- Fibonacci Heaps

Run-time Per Operation

	Insert	DeleteMin	Merge ($=H_1+H_2$)
 Binary heap	<ul style="list-style-type: none"> ■ $O(\log n)$ worst-case ■ $O(1)$ amortized for buildHeap 	$O(\log n)$	$O(n)$
Leftist Heap	$O(\log n)$	$O(\log n)$	$O(\log n)$
Skew Heap	$O(\log n)$	$O(\log n)$	$O(\log n)$
 Binomial Heap	<ul style="list-style-type: none"> ■ $O(\log n)$ worst case ■ $O(1)$ amortized for sequence of n inserts 	$O(\log n)$	$O(\log n)$
Fibonacci Heap	$O(1)$	$O(\log n)$	$O(1)$

Priority Queues in STL

- Uses *Binary heap*
- Default is MaxHeap
- Methods
 - Push, top, pop, empty, clear


```
#include <priority_queue>
int main ()
{
    priority_queue<int> Q;
    Q.push (10);
    cout << Q.top ();
    Q.pop ();
}
```

Calls DeleteMax()

For MinHeap: declare priority_queue as:

```
priority_queue<int, vector<int>, greater<int>> Q;
```

Refer to Book Chapter 6, Fig 6.57 for an example



Binomial Heaps



Binomial Heap

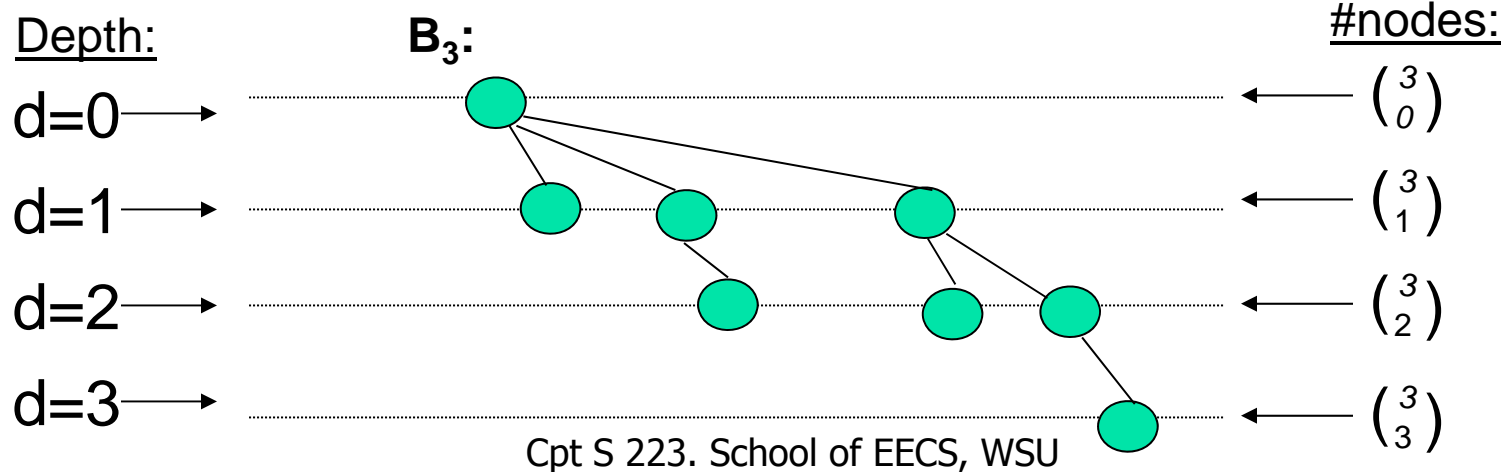
- A binomial heap is a *forest* of heap-ordered binomial trees, satisfying:
 - i) Structure property, and
 - ii) Heap order property
- A binomial heap is different from binary heap in that:
 - Its structure property is totally different
 - Its heap-order property (within each binomial tree) is the same as in a binary heap

Note: A binomial tree *need not* be a binary tree

Definition: A "Binomial Tree" B_k

- A binomial tree of height k is called B_k :
 - It has 2^k nodes
 - The number of nodes at depth $d = \binom{k}{d}$

$\binom{k}{d}$ is the form of the co-efficients in binomial theorem





What will a Binomial Heap with $n=31$ nodes look like?

- We know that:
 - i) A binomial heap should be a forest of binomial trees
 - ii) Each binomial tree has power of 2 elements
- So how many binomial trees do we need?

$$n = 31 = \begin{matrix} B_4 & B_3 & B_2 & B_1 & B_0 \\ (1 & 1 & 1 & 1 & 1)_2 \end{matrix}$$

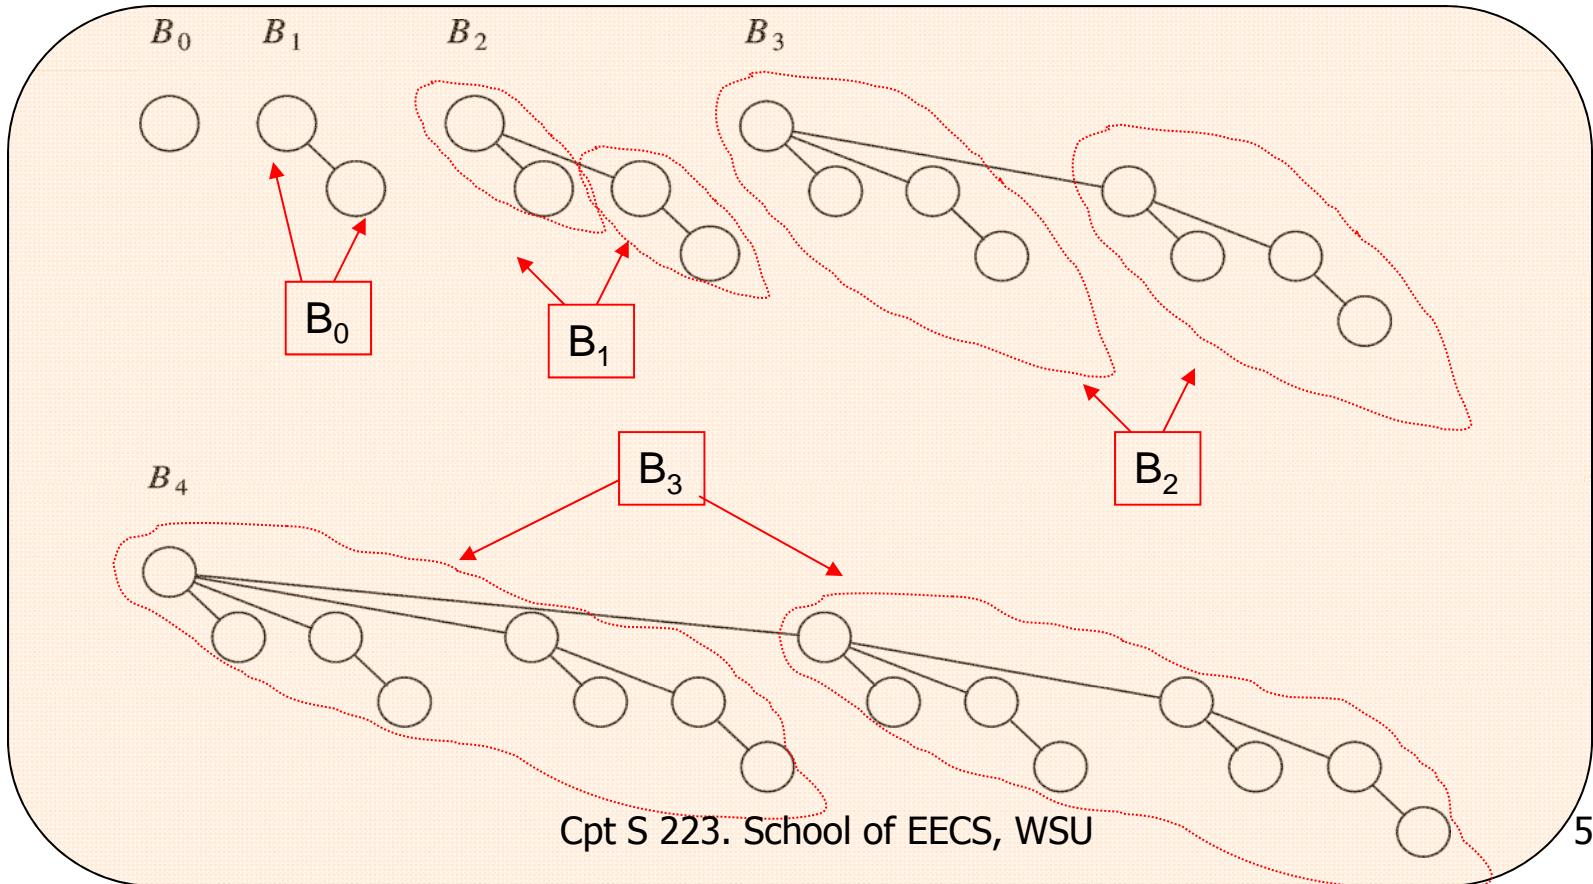
A Binomial Heap w/ $n=31$ nodes

$B_4 B_3 B_2 B_1 B_0$

$$n = 31 = (1\ 1\ 1\ 1\ 1)_2$$

$$B_i == B_{i-1} + B_{i-1}$$

Forest of binomial trees $\{B_0, B_1, B_2, B_3, B_4\}$





Binomial Heap Property

- Lemma: There exists a binomial heap for every positive value of n
- Proof:
 - All values of n can be represented in binary representation
 - Have one binomial tree for each power of two with co-efficient of 1
 - Eg., $n=10 \implies (1010)_2 \implies$ forest contains $\{B_3, B_1\}$



Binomial Heaps: *Heap-Order Property*

- Each binomial tree should contain the minimum element at the root of every subtree
 - Just like binary heap, except that the tree here is a binomial tree structure (and not a complete binary tree)
- The order of elements across binomial trees is irrelevant





Definition: Binomial Heaps

- A binomial heap of n nodes is:
 - (**Structure Property**) A forest of binomial trees as dictated by the binary representation of n
 - (**Heap-Order Property**) Each binomial tree is a min-heap or a max-heap

Binomial Heaps: Examples

Two different heaps:





Key Properties

- Could there be multiple trees of the same height in a binomial heap?

no

- What is the upper bound on the number of binomial trees in a binomial heap of n nodes?

$\lceil \lg n \rceil$


- Given n , can we tell (for sure) if B_k exists?

B_k exists if and only if:
the k^{th} least significant bit is 1
in the binary representation of n

An Implementation of a Binomial Heap

Maintain a linked list of tree pointers (for the forest)

Example: $n=13 == (1101)_2$



Shown using the left-child, right-sibling pointer method

Analogous to a bit-based representation of a binary number n



Binomial Heap: Operations

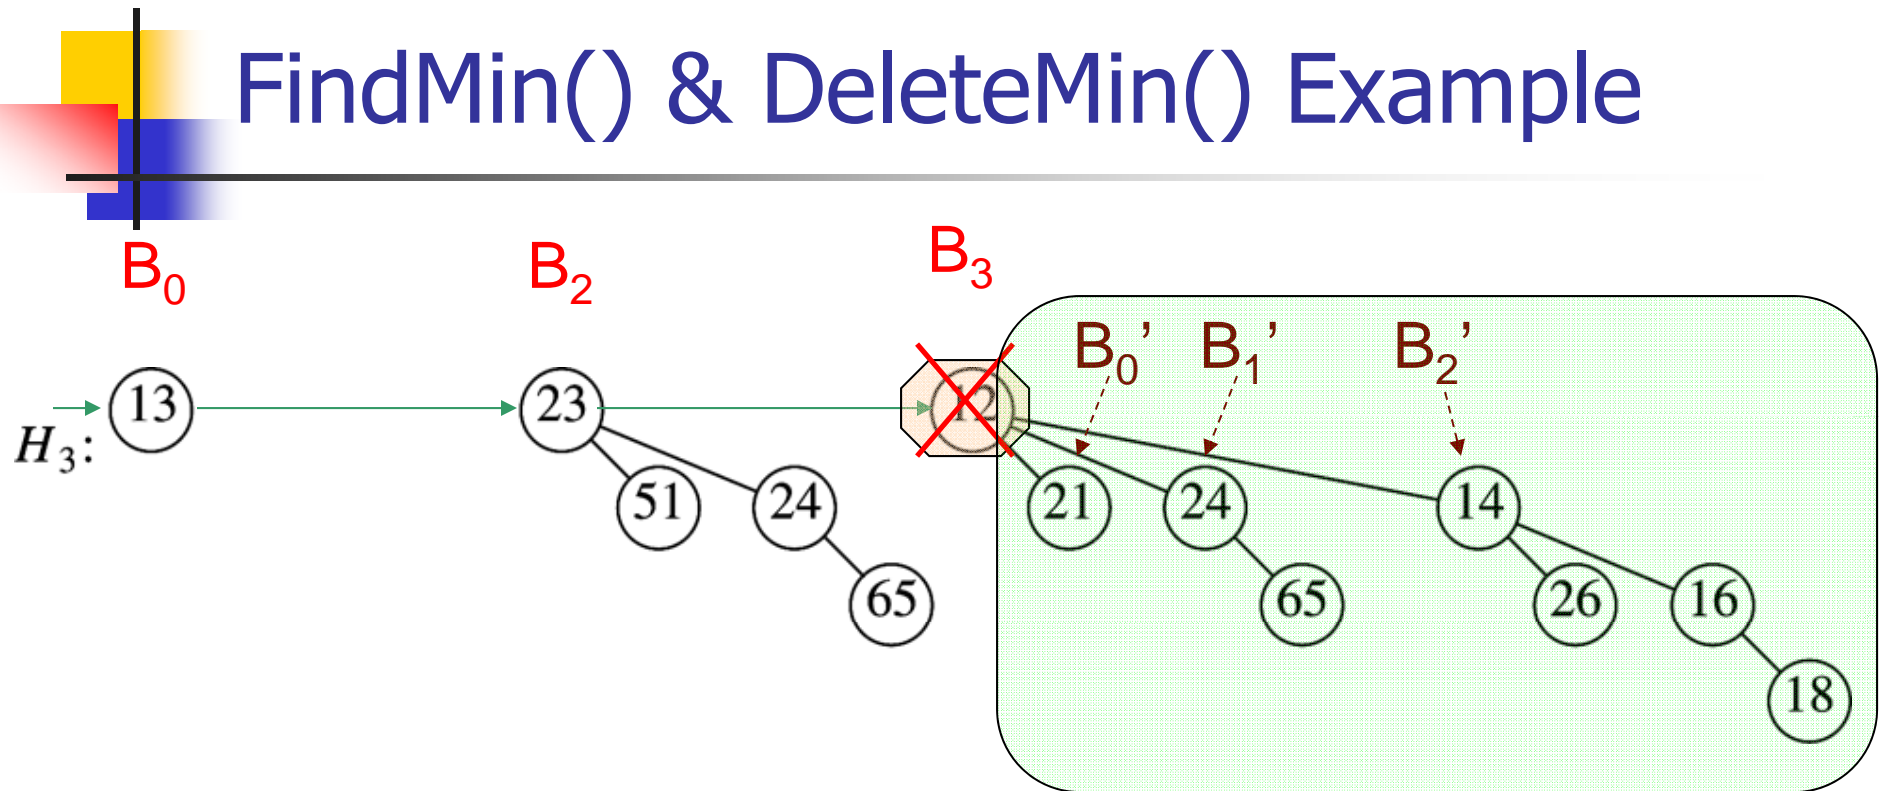
- $x \leq \text{DeleteMin}()$
- $\text{Insert}(x)$
- $\text{Merge}(H_1, H_2)$



DeleteMin()

- **Goal:** Given a binomial heap, H , find the minimum and delete it
- **Observation:** The root of each binomial tree in H contains its minimum element
- **Approach:** Therefore, return the minimum of all the roots (minimums)
- **Complexity:** $O(\log n)$ comparisons
(because there are only $O(\log n)$ trees)

FindMin() & DeleteMin() Example



For DeleteMin(): After delete, how to adjust the heap?

New Heap : Merge $\{ B_0, B_2 \}$ & $\{ B_0', B_1', B_2' \}$



Insert(x) in Binomial Heap

- **Goal:** To insert a new element x into a binomial heap H
- **Observation:**
 - Element x can be viewed as a single element binomial heap
 - \Rightarrow Insert (H, x) == **Merge($H, \{x\}$)**

So, if we decide how to do merge we will automatically figure out how to implement both insert() and deleteMin()

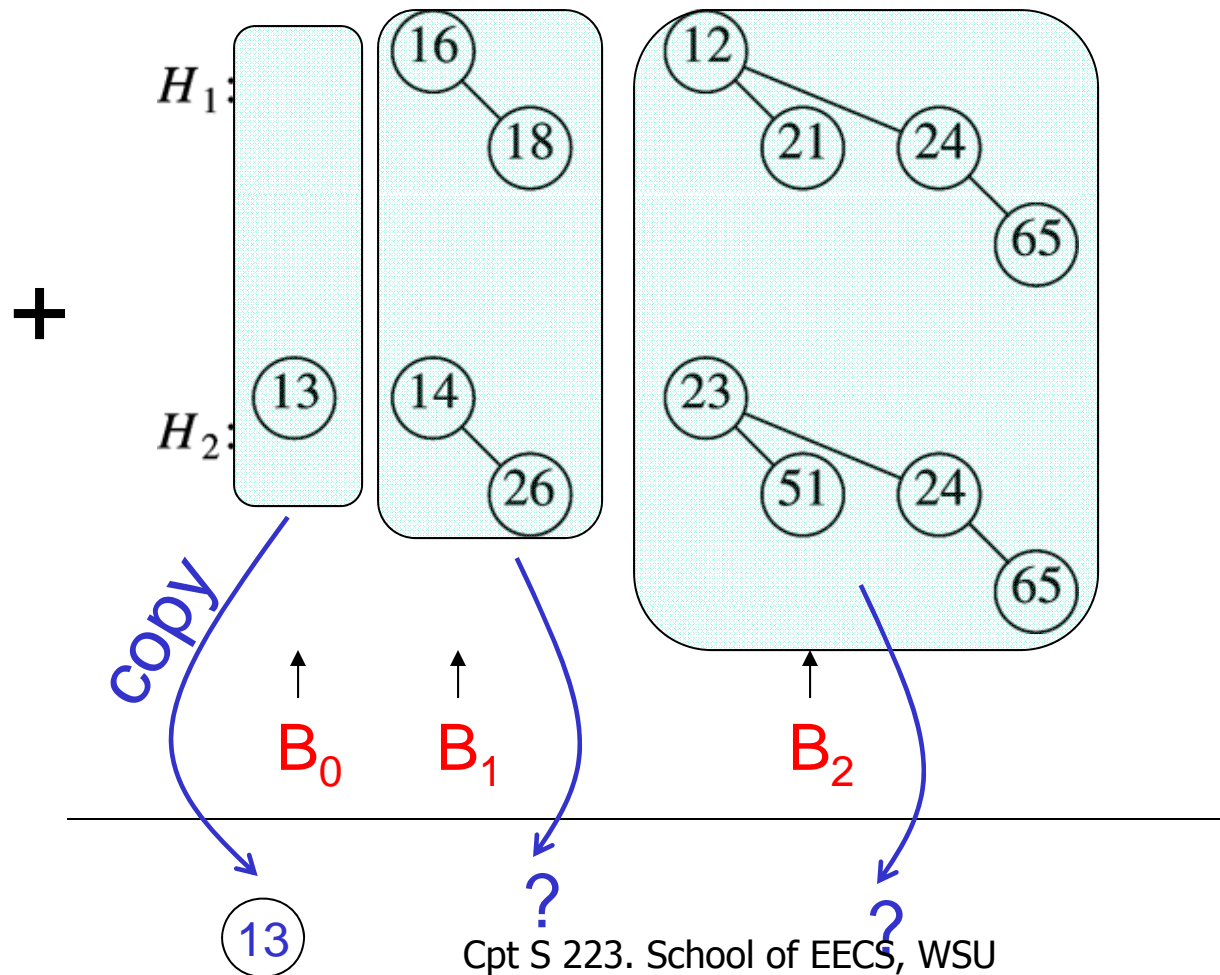


Merge(H_1, H_2)


- Let n_1 be the number of nodes in H_1
- Let n_2 be the number of nodes in H_2
- Therefore, the new heap is going to have $n_1 + n_2$ nodes
 - Assume $n = n_1 + n_2$
- Logic:
 - Merge trees of same height, starting from lowest height trees
 - If only one tree of a given height, then just copy that
 - Otherwise, need to do carryover (just like adding two binary numbers)

Idea: merge tree of same heights

Merge: Example



How to Merge Two Binomial Trees of the *Same* Height?

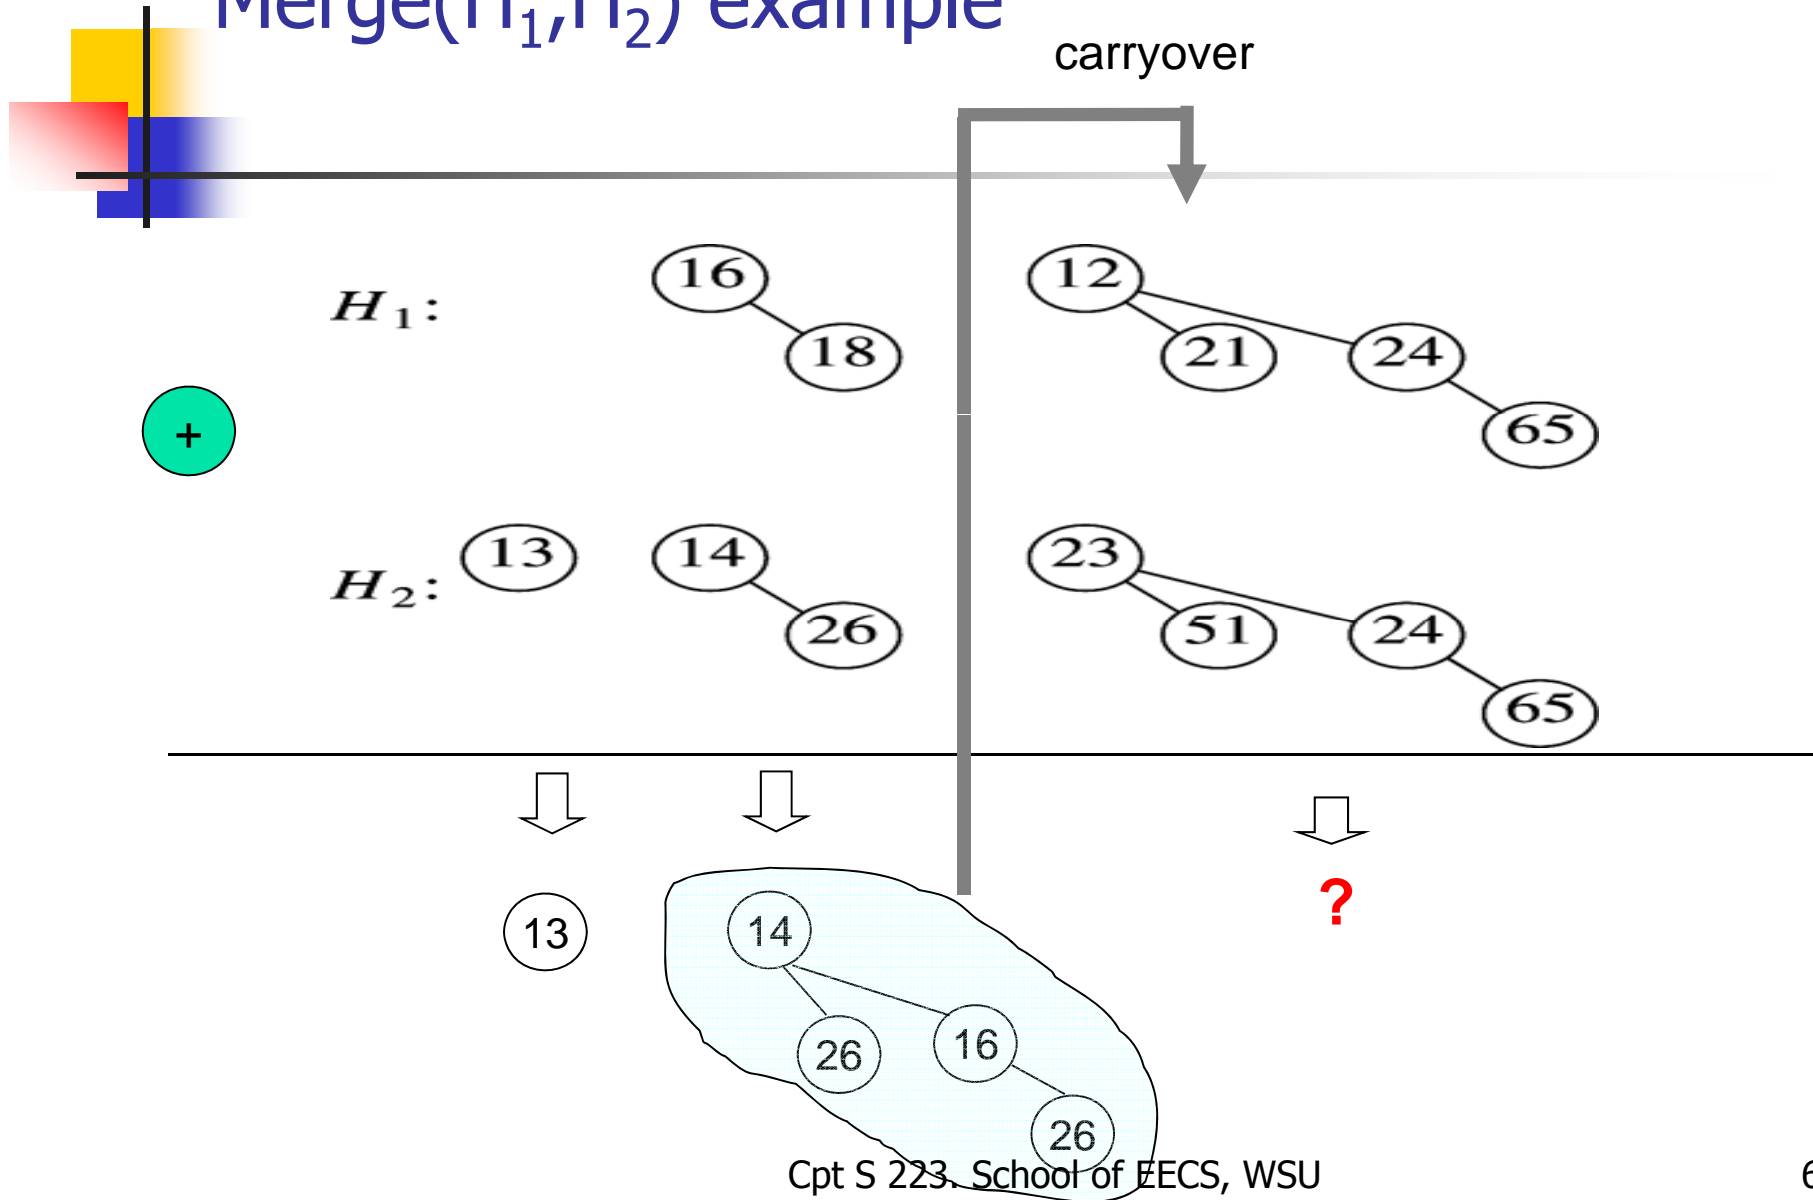


Simply compare the roots

```
1  /**
2   * Return the result of merging equal-sized t1 and t2.
3   */
4  BinomialNode * combineTrees( BinomialNode *t1, BinomialNode *t2 )
5  {
6      if( t2->element < t1->element )
7          return combineTrees( t2, t1 );
8      t2->nextSibling = t1->leftChild;
9      t1->leftChild = t2;
10     return t1;
11 }
```


Note: Merge is defined for only binomial trees with the same height

Merge(H_1, H_2) example




How to Merge *more than two* binomial trees of the same height?

- Merging more than 2 binomial trees of the same height could generate carry-overs



Merge any two and leave the third as carry-over




Merge(H_1, H_2) : Example

Input:

H_1 :


+

H_2 :



Output:

H_3 :




There are two other possible answers

Merge cost $\propto \log(\max\{n_1, n_2\}) = O(\log n)$ comparisons

Run-time Complexities

- Merge takes $O(\log n)$ comparisons
- Corollary:
 - Insert and DeleteMin also take $O(\log n)$
- *It can be further proved that an uninterrupted sequence of m Insert operations takes only $O(m)$ time per operation, implying $O(1)$ amortize time per insert*
 - Proof Hint:
 - For each insertion, if i is the least significant bit position with a 0, then number of comparisons required to do the next insert is $i+1$
 - If you count the #bit flips for each insert, going from insert of the first element to the insert of the last (n^{th}) element, then
=> amortized run-time of $O(1)$ per insert

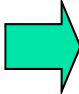
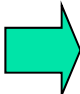




Binomial Queue Run-time Summary

- Insert
 - $O(\lg n)$ worst-case
 - $O(1)$ amortized time if insertion is done in an uninterrupted sequence (i.e., without being intervened by deleteMins)
- DeleteMin, FindMin
 - $O(\lg n)$ worst-case
- Merge
 - $O(\lg n)$ worst-case

Run-time Per Operation

	Insert	DeleteMin	Merge ($=H_1+H_2$)
 Binary heap	<ul style="list-style-type: none"> $O(\log n)$ worst-case $O(1)$ amortized for buildHeap 	$O(\log n)$	$O(n)$
Leftist Heap	$O(\log n)$	$O(\log n)$	$O(\log n)$
Skew Heap	$O(\log n)$	$O(\log n)$	$O(\log n)$
 Binomial Heap	<ul style="list-style-type: none"> $O(\log n)$ worst case $O(1)$ amortized for sequence of n inserts 	$O(\log n)$	$O(\log n)$
Fibonacci Heap	$O(1)$	$O(\log n)$	$O(1)$



Summary

- Priority queues maintain the minimum or maximum element of a set
- Support $O(\log N)$ operations worst-case
 - insert, deleteMin, merge
- Many applications in support of other algorithms