

Binary Heaps

CSE 373 - Data Structures

April 26, 2002

Readings and References

- Reading
 - › Sections 6.1-6.4, *Data Structures and Algorithm Analysis in C*, Weiss
- Other References

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

2

A New Problem...

- Application: Find the smallest (or highest priority) item quickly
 - › Operating system needs to schedule jobs according to priority
 - › Doctors in ER take patients according to severity of injuries
 - › Event simulation (bank customers arriving and departing, ordered according to when the event happened)

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

3

Use Lists or Binary Search Tree?

- We want an ADT that can efficiently perform:
 - › FindMin (and DeleteMin)
 - › Insert
- What if we use...
 - › Lists: If sorted, what is the run time for Insert and FindMin? Unsorted?
 - › Binary Search Trees: What is the run time for Insert and FindMin?

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

4

Less flexibility → More speed

- Lists
 - › If sorted: FindMin is $O(1)$ but Insert is $O(N)$
 - › If not sorted: Insert is $O(1)$ but FindMin is $O(N)$
- Binary Search Trees (BSTs)
 - › Insert is $O(\log N)$ and FindMin is $O(\log N)$
- BSTs look good but...
 - › BSTs are efficient for all Finds, not just FindMin
 - › We only need FindMin

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

5

Better than a speeding BST

- We can do better than Binary Search Trees
 - › Very limited requirements: Insert, FindMin, DeleteMin
 - › FindMin is $O(1)$
 - › Insert is $O(\log N)$
 - › DeleteMin is $O(\log N)$

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

6

Binary Heaps

- A binary heap is a binary tree that is:
 - › Complete: the tree is completely filled except possibly the bottom level, which is filled from left to right
 - › Satisfies the heap order property
 - every node is less than or equal to its children
 - or every node is greater than or equal to its children
- The root node is always the smallest node
 - › or the largest, depending on the heap order




26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

7

Heap order property

- A heap provides limited ordering information
- Each *path* is sorted, but the subtrees are not sorted relative to each other
 - › A binary heap is NOT a binary search tree



These are all valid binary heaps (minimum)


26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

8


Binary Heap vs Binary Search Tree

Binary Heap



Parent is less than both left and right children

Binary Search Tree

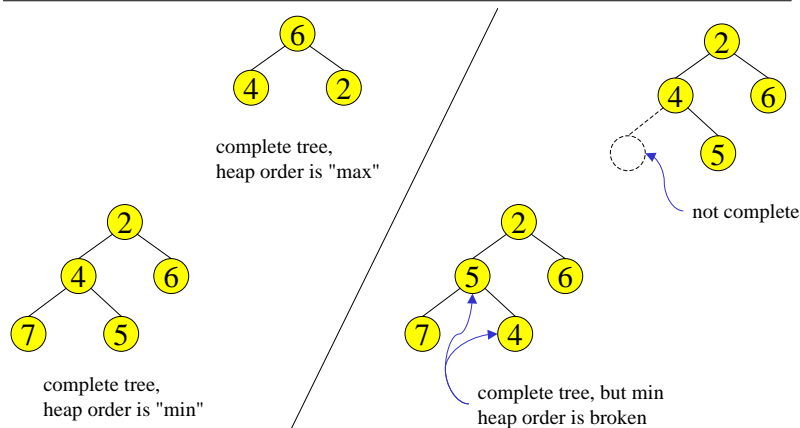


Parent is greater than left child, less than right child

Structure property


- A binary heap is a complete tree
 - › All nodes are in use except for possibly the right end of the bottom row
- Pointers from node to node?
 - › allow arbitrary connect and disconnect at any node
 - › but we don't need this flexibility since the tree is always complete and we don't need to do a lot of reorganizing to meet a tree order property

Examples



Array Implementation of Heaps


- Root node = $A[1]$
- Children of $A[i] = A[2i], A[2i + 1]$
- Keep track of current size N (number of nodes)



FindMin and DeleteMin

- FindMin: Easy!

- › Return root value A[1]
- › Run time = ?




- DeleteMin:

- › Delete (and return) value at root node

DeleteMin


- Delete (and return) value at root node




Maintain the Structure Property

- We now have a “Hole” at the root

- › Need to fill the hole with another value




- When we get done, the tree will have one less node and must still be complete




Maintain the Heap Property

- The last value has lost its node
 - › we need to find a new place for it
- We can do a simple insertion sort operation to find the correct place for it in the tree



DeleteMin: Percolate Down



- Keep comparing with children $A[2i]$ and $A[2i + 1]$
- Copy smaller child up and go down one level
- Done if both children are \geq item or reached a leaf node
- What is the run time?

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

17

DeleteMin: Run Time Analysis

- Run time is $O(\text{depth of heap})$
- A heap is a complete binary tree
- Depth of a complete binary tree of N nodes?
 - › $\text{depth} = \lfloor \log(N) \rfloor = \text{floor}(\log(N))$
- Run time of DeleteMin is $O(\log N)$


26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

18

Insert

- Add a value to the tree
- Structure and heap order properties must still be correct when we are done




26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

19

Maintain the Structure Property

- The only valid place for a new node in a complete tree is at the end of the array
- We need to decide on the correct value for the new node, and adjust the heap accordingly




26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps


20

Maintain the Heap Property

- The new value goes where?
- We can do a simple insertion sort operation to find the correct place for it in the tree




Insert: Percolate Up



- Start at last node and keep comparing with parent $A[i/2]$
- If parent larger, copy parent down and go up one level
- Done if parent \leq item or reached top node $A[1]$
- Run time?


Insert: Done



- Run time?

Sentinel Values

- Every iteration of Insert needs to test:
 - > if it has reached the top node $A[1]$
 - > if parent \leq item
- Can avoid first test if $A[0]$ contains a very large negative value
 - > sentinel $-\infty <$ item, for all items
- Second test alone always stops at top



value	$-\infty$	2	3	8	7	4	10	9	11	9	6	5		
index	0	1	2	3	4	5	6	7	8	9	10	11	12	13

Summary of Heap ADT Analysis

- Space needed for heap of N nodes: $O(\text{MaxN})$
 - › An array of size MaxN , plus a variable to store the size N , plus an array slot to hold the sentinel
- Time
 - › FindMin: $O(1)$
 - › DeleteMin and Insert: $O(\log N)$
 - › BuildHeap from N inputs
 - N Insert operations = $O(N \log N)$
 - Treat input array as a heap and fix it using percolate down = $O(N)$

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

25

Other Heap Operations

- Find(X, H): Find the element X in heap H of N elements
 - › What is the running time? $O(N)$
- FindMax(H): Find the maximum element in H
 - › What is the running time? $O(N)$
- We sacrificed performance of these operations in order to get $O(1)$ performance for FindMin

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

26

Other Heap Operations

- DecreaseKey(P, Δ, H): Decrease the key value of node at position P by a positive amount Δ . eg, to increase priority
 - › First, subtract Δ from current value at P
 - › Heap order property may be violated
 - › so percolate up to fix
 - › Running Time: $O(\log N)$

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

27

Other Heap Operations

- IncreaseKey(P, Δ, H): Increase the key value of node at position P by a positive amount Δ . eg, to decrease priority
 - › First, add Δ to current value at P
 - › Heap order property may be violated
 - › so percolate down to fix
 - › Running Time: $O(\log N)$

26-Apr-02

CSE 373 - Data Structures - 11 - Binary Heaps

28

Other Heap Operations

- Delete(P,H): E.g. Delete a job waiting in queue that has been preemptively terminated by user
 - › Use DecreaseKey(P, ∞ ,H) followed by DeleteMin
 - › Be careful about your sentinel value and overflow
 - › Running Time: $O(\log N)$

Other Heap Operations

- Merge(H1,H2): Merge two heaps H1 and H2 of size $O(N)$. H1 and H2 are stored in two arrays.
 - › Can do $O(N)$ Insert operations: $O(N \log N)$ time
 - › Better: Copy H2 at the end of H1 and use BuildHeap. Running Time: $O(N)$
- Merges in $O(\log N)$ coming soon to a lecture hall near you ...