CSCI-1200 Data Structures — Spring 2021
Lecture 23 – Priority Queues

Review from Lectures 20 & 21

- Hash Tables, Hash Functions, and Collision Resolution

- Performance of: Hash Tables vs. Binary Search Trees

- STL's `unordered_set` (and `unordered_map`)

- Using a hash table to implement a set/map
  - Hash functions as functors/function objects
  - Iterators, find, insert, and erase

- Using STL's `for_each`

- STL Queue and STL Stack

- What's a Priority Queue?

Today's Lecture

- A Priority Queue as a Heap

- A Heap as a Vector

- Building a Heap

- Heap Sort

23.1   Implementing Pop (a.k.a. Delete Min)

- The value at the top (root) of the tree is replaced by the value stored in the last leaf node.
  This has echoes of the erase function in binary search trees.

- The last leaf node is removed.
  QUESTION: But how do we find the last leaf? Ignore this for now...

- The value now at the root likely breaks the heap property. We use the `percolate_down` function to restore the heap property. This function is written here in terms of tree nodes with child pointers (and the priority stored as a `value`), but later it will be written in terms of vector subscripts.

```
percolate_down(TreeNode<T> * p) {
  while (p->left) {
    TreeNode<T>* child;
    //  Choose the child to compare against
    if (p->right && p->right->value < p->left->value)
      child = p->right;
    else
      child = p->left;
    if (child->value < p->value) {
      swap(child, p); // value and other non-pointer member vars
      p = child;
    }
    else
      break;
  }
}
```

23.2   Implementing Push (a.k.a. Insert)

- To add a value to the heap, a new last leaf node in the tree is created to store that value.

- Then the `percolate_up` function is run. It assumes each node has a pointer to its parent.

```
percolate_up(TreeNode<T> * p) {
  while (p->parent)
    if (p->value < p->parent->value) {
      swap(p, parent);  // value and other non-pointer member vars
      p = p->parent;
    }
    else
      break;
}
```

23.3   Push (Insert) and Pop (Delete-Min) Usage Exercise

Suppose the following operations are applied to an initially empty binary heap of integers. Show the resulting heap after each `delete_min` operation. (Remember, the tree must be **complete**!)

```
    push 5, push 3, push 8, push 10, push 1, push 6,
    pop,
    push 14, push 2, push 4, push 7,
    pop,
    pop,
    pop
```

23.4   Heap Operations Analysis

- Both `percolate_down` and `percolate_up` are $O(\log n)$ in the worst-case. Why?

- But, `percolate_up` (and as a result `push`) is $O(1)$ in the average case. Why?

23.5   Implementing a Heap with a Vector (instead of Nodes & Pointers)

- In the vector implementation, the tree is never explicitly constructed. Instead the heap is stored as a vector, and the child and parent "pointers" can be implicitly calculated.

- To do this, number the nodes in the tree starting with 0 first by level (top to bottom) and then scanning across each row (left to right). These are the vector indices. Place the values in a vector in this order.

- As a result, for each subscript, i,
   - The parent, if it exists, is at location $\lfloor (i-1)/2 \rfloor$.
   - The left child, if it exists, is at location $2i + 1$.
   - The right child, if it exists, is at location $2i + 2$.

- For a binary heap containing n values, the last leaf is at location $n - 1$ in the vector and the first node with less than two children is at location $\lfloor (n-1)/2 \rfloor$.

- The standard library (STL) `priority_queue` is implemented as a binary heap.

23.6   Heap as a Vector Exercises

- Draw a binary heap with values: 52 13 48 7 32 40 18 25 4, first as a tree of nodes & pointers, then in vector representation.

- Starting with an initially empty heap, show the vector contents for the binary heap after each `delete_min` operation.

```
push 8, push 12, push 7, push 5, push 17, push 1,
pop,
push 6, push 22, push 14, push 9,
pop,
pop,
```

23.7   Building A Heap

- In order to build a heap from a vector of values, for each index from $\lfloor (n-1)/2 \rfloor$ down to 0, run `percolate_down`. Show that this fully organizes the data as a heap and requires at most $O(n)$ operations.

- If instead, we ran `percolate_up` from each index starting at index 0 through index n-1, we would get properly organized heap data, but incur a $O(n \log n)$ cost. Why?

23.8   Heap Sort

- Heap Sort is a simple algorithm to sort a vector of values: Build a heap and then run n consecutive `pop` operations, storing each "popped" value in a new vector.

- It is straightforward to show that this requires $O(n \log n)$ time.

- **Exercise:**   Implement an *in-place* heap sort. An in-place algorithm uses only the memory holding the input data – a separate large temporary vector is not needed.

23.9   Summary Notes about Vector-Based Priority Queues

- Priority queues are conceptually similar to queues, but the order in which values / entries are removed ("popped") depends on a priority.

- Heaps, which are conceptually a binary tree but are implemented in a vector, are the data structure of choice for a priority queue.

- In some applications, the priority of an entry may change while the entry is in the priority queue. This requires that there be "hooks" (usually in the form of indices) into the internal structure of the priority queue. This is an implementation detail we have not discussed.