

Priority Queues & Heaps

- A *priority queue* is simply like a queue except that the next element to be dequeued always has the maximum priority.
- A priority queue is an *abstract data type*, which means that we don't know what data structure lies beneath it. All we know is the *interface* i.e. we have a way to insert into it (`pq_insert`) and to get the element of maximum priority from it (`pq_delmin`).
- We usually think of the element with the least integer value being the element with maximum priority. Hence `pq_delmin` gives the element with maximum priority. Also, this means our heaps are *min-heaps*, in contrast to *max-heaps* in which the elements with larger integer value have larger priority.
- A heap is an *implementation* for a priority queue.
- A heap has two important invariants:
 - (a) **Ordering Invariant** : Any element in the heap is less than or equal to its children.
 - (b) **Shape invariant**: An element is inserted in a manner that, at any point, only the last level is unfilled and elements are filled in this level from left to right.
- As mentioned in lecture, for now, we'll think of an array that stores the values in the heap, such that the children of the element at index i are at $2i$ and $2i + 1$
Notice that this means that we can't use the index 0 of the array!
- While writing code, we'll see that we need one important invariant for heaps — the element which will be accessed/deleted from the heap must have minimum value. To maintain this invariant, while inserting or deleting from the priority queue, we need to do some special operations.
- We'll also see some important concepts that are very common while dealing with operations on complex data structures:
 - (a) The invariants that hold for the operation as a whole (which are true at the beginning and at the end of the operation) are not necessarily true while the operation is being performed. So we are permitted to **temporarily violate invariants** while performing the operation and then restore them at the end.
 - (b) Thus, to prove correctness, we need to find variations of the original invariants that hold true while the operation is being performed and reason that these will imply the stronger invariants that are true at the end of the operation.

Checkpoint 0

When representing heaps as arrays, why can't we use the index 0 of the array?

Playing with heaps!

Use the visualization at <http://www.cs.usfca.edu/~galles/visualization/Heap.html> to insert the following elements into the heap, in the given order.

5, 3, 6, 7, 2, 6

(Here's some space if you want to draw the final heap obtained)

Some observations about the heap you see now:

- The smallest element in the heap (i.e. 2) is at the root. Convince yourself that this must always hold true because of the ordering invariant
- The largest element in the heap (i.e. 7) is in the last level of the heap. Again, convince yourself that this must always hold true by the ordering invariant.
- The heap can contain duplicates, since the ordering invariant states that the children of an element must be greater than *or equal to* the element

Now perform a few operations of removing the smallest element from the heap.

Checkpoint 1

What is the time complexity of the insert and remove minimum functions? Justify your answer.

Discussing the code for deletion from a heap

When we delete the minimum element from the heap, we notice that while we're performing the operation, the ordering invariant of the heap is violated i.e. the intermediate heaps don't satisfy the ordering invariant.

To prove that the ordering invariant is satisfied at the end of the delete operation, we need to come up with an invariant that will hold of the intermediate heaps and that is strong enough to prove the ordering invariant for the final heap.

Turns out an invariant that holds at each step is that the heap invariant is satisfied at all points, except at the current element, which may be larger than its children. We express this invariant as the `is_heap_except_down` function.

```

1 bool is_heap_except_down(heap H, int n) {
2   if (!is_safe_heap(H)) return false;
3
4   /* check parent <= node for all nodes except root (i = 1) */
5   /* and children of n (i/2 = n) */
6   for (int i = 2; i < H->next; i++)
7     //@loop_invariant 2 <= i;
8     {
9       if (i/2 != n && !(priority(H, i/2) <= priority(H, i))) {
10        return false;
11      }
12      /* for children of node n, check grandparent */
13      if (i/2 == n && (i/2)/2 >= 1 && !(priority(H, (i/2)/2) <= priority(H, i))) {
14        return false;
15      }
16    }
17   return true;
18 }

```

Using this `is_heap_except_down` function as an invariant, we can write the `sift_down` and `pq_delmin` function as follows:

```

1 void sift_down(heap H, int i)
2 //@requires 1 <= i && i < H->next;
3 //@requires is_heap_except_down(H, i);
4 //@ensures is_heap(H);
5 {
6   int n = H->next;
7   int left = 2*i;
8   int right = left + 1;
9   while (left < n)
10    //@loop_invariant 1 <= i && i < n;
11    //@loop_invariant left == 2*i && right == 2*i+1;
12    //@loop_invariant is_heap_except_down(H, i);
13    {
14      if (priority(H,i) <= priority(H,left)
15          && (right >= n || priority(H,i) <= priority(H,right))) {
16        return;
17      }
18      if (right >= n || priority(H,left) < priority(H,right)) {
19        swap(H->data, i, left);
20        i = left;
21      }
22      else {
23        //@assert right < n && priority(H, right) <= priority(H,left);
24        swap(H->data, i, right);
25        i = right;
26      }
27      left = 2*i;
28      right = left+1;
29    }
30    //@assert i < n && 2*i >= n;
31    //@assert is_heap_except_down(H, i);
32    return;
33 }
34

```

```
35 elem pq_delmin(heap H)
36 //@requires is_heap(H) && !pq_empty(H);
37 //@ensures is_heap(H);
38 {
39     int n = H->n;
40     elem min = H->data[1];
41     H->data[1] = H->data[n-1];
42     H->n = n-1;
43     if (H->n > 1) sift_down(H, 1);
44     return min;
45 }
```

Checkpoint 2

(Optional) Come up with a proof of correctness for the `pq_delmin` function. This need not be super formal, but you should be able to convince yourself that the loop invariant is preserved and that the loop invariant and negation of the loop guard imply the postcondition.