

HW4: Binary Heaps


Due: Thursday, November 2, at 11:59 PM, on GSC

You may work on your own or with one (1) partner.


In this assignment, you will implement a fixed-size binary heap. The structure of the heap is already defined for you in `binheap.rkt`¹. The heap is represented using a DSSL vector to contain the elements. Each heap also contains a comparison function for ordering the elements of the heap, so that your implementation can support heaps of integers, heaps of strings, heaps of whatits, heaps of sporkles, etc.

Representation overview

A binary heap is a *complete* binary tree satisfying the *heap property*. Completeness means that each level is full except the last one, and the heap property means that each node's value is less than its children's values. Here is an example of a binary heap drawn as a tree:



In practice, binary trees are rarely represented as actual tree nodes and links. Instead, the completeness property means that the elements can be stored in level order in an array, which is more space- and cache-efficient. In particular, in this assignment we will represent a heap as a struct of three fields. Here is the same heap as above in our actual heap representation:



The third field, `data`, stores a reference to the vector holding the elements. In order to allow the heap to grow, we do not necessarily use the full capacity

¹<http://goo.gl/5v3p4R>

of the vector; instead we store the actual number of elements in `size`, the first field.

The second field, `lt?`, stores the *heap ordering predicate*. This is necessary because your heap must work with any element type whatsoever, not just with numbers or strings. If elements were only numbers or strings, you could compare them with `x, y: x < y`, but that won't work for, say, bank account structs, employees, or weighted graph edges. So each `HeapOf [X]` has a field `lt?` containing the correct function for comparing its elements. Supposing that `h` is a `HeapOf [X]`, and `a` and `b` are two elements (*i.e.*, `X`s), we can compare them with `h.lt?(a, b)`, which answers `True` when `a` is less than `b` and `False` otherwise.

Your task

In `binheap.rkt`, I've supplied a definition of a function `create_heap` that returns a new, empty heap given a capacity and an ordering function. Implementing the remaining operations is up to you:

```
heap_insert! : HeapOf [X] X -> Void      ;  $\mathcal{O}(\log n)$ 
find_min     : HeapOf [X] -> X           ;  $\mathcal{O}(1)$ 
remove_min!  : HeapOf [X] -> Void      ;  $\mathcal{O}(\log n)$ 
```

For details, see the function headers provided in `binheap.rkt`, which include purpose statements as well. Each operation must have the worst-case shown above, where n is the number of elements in the heap.

In order to help you factor your program effectively, I've included at the bottom of `binheap.rkt` a list of helper functions with names, signatures (types), and purpose statements (brief functional descriptions). You are free to use as much or as little of my design as you like.

Brownie points (not real points)

Make your heap expand as necessary to accommodate any number of insertions. To achieve this, instead of failing when the heap is full, `heap_insert!` should allocate a new vector that doubles the capacity and copy the existing elements over from the old vector.

Deliverables

- The provided file `binheap.rkt`, containing:
 - the `heap_insert!`, `find_min`, and `remove_min!` functions fully defined, and
 - sufficient tests to convince yourself your code's correctness.