Assignment 5: Priority Queue

Assignment handout authors include: Cynthia Lee, Marty Stepp, and Julie Zelenski.
Some edits over time by Jerry Cain, Keith Schwarz, and others.

Due: Wednesday, November 5[th] at 5:00pm

Pair programming <u>is</u> permitted on this assignment. See course information sheet and honor code.

On the very first day of class, I talked about the structure of our course overall as follows: in the beginning, we learn about the ADTs from the client perspective, then (after an aside to study recursion) we learn about how the ADTs could be implemented behind the scenes in C++. We are now firmly in the "behind the scenes" stage of the course, with lectures on ArrayList implementation, Priority Queue implementation, and Map implementation. For this assignment, you'll be operating directly in the same space by implementing a PriorityQueue in several different ways. It's the most low-level assignment of the quarter, which can be challenging, but it's an irreplaceable learning experience in your programming life. ☺

The assignment has four main purposes:
1. To give you experience implementing a data structure "behind the scenes."
2. To give you experience working with pointers and dynamic memory.
3. To give you experience working with several classic data structures.
4. To compare, using empirical experiments, the performance of a data structures that share the same interface but have different implementations.

Deliverables (turn in these files):
- `SLinkedPriorityQueue.h/.cpp:` a PQ implementation using a <u>sorted linked list</u> as internal storage
- `USLinkedPriorityQueue.h/.cpp:` a PQ implementation using an <u>unsorted linked list</u> as internal storage
- `HeapPriorityQueue.h/.cpp:` a PQ implementation using a <u>"heap"</u> array as internal storage
- `analysis.pdf:` a written document reporting your results on empirical experiments with your code's performance (you may use any program to do the writeup, but please <u>only turn in PDF</u>, not .doc/.odt/etc.)
- `pqueue-main.cpp:` a simple test harness to run tests necessary for your analysis.pdf analysis.

This is a pair assignment. If you work as a pair, comment both members' names on top of every submitted code file. Only one of you should submit the assignment; do not turn in two copies.

THE PRIORITY QUEUE COLLECTION

A priority queue is a collection that is similar to a queue, except that the elements are enqueued with "priority" ratings. Then when the elements are dequeued later, they always come out in increasing order of priority. That is, regardless of the order in which you add (enqueue) the elements, when you remove (dequeue) them, the one with the lowest priority number comes out first, then the second-smallest, and so on, with the largest priority item coming out last. Priority queues are useful for modeling tasks where some jobs are more important to process than others, like the line of patients at an emergency room (a severely injured patient has higher "priority" than one with a cold).

We will use the convention that a **<u>smaller priority number means a greater urgency</u>**, such that a priority-1 item would take precedence over a priority-2 item. Because terms like "higher priority" can be confusing, since a "higher" priority means a lower integer value, we will follow the convention in this document of referring to a "more urgent" priority for a smaller integer and a "less urgent" priority for a greater integer.

In this assignment, you will be writing three different implementations of a priority queue class that stores strings. If two strings in the queue have the same priority, you will break ties by considering the one that comes first in alphabetical order to come first. Use C++'s built-in relational operators (`<`, `>`, `<=`, etc.) to compare the strings.

For example, if you enqueue these strings into a priority queue:
> enqueue **"x"** with priority 5
> enqueue **"b"** with priority 4
> enqueue **"a"** with priority 8
> enqueue **"m"** with priority 5
> enqueue **"q"** with priority 5
> enqueue **"t"** with priority 2

Then if you were to dequeue the strings, they would be returned in this order:
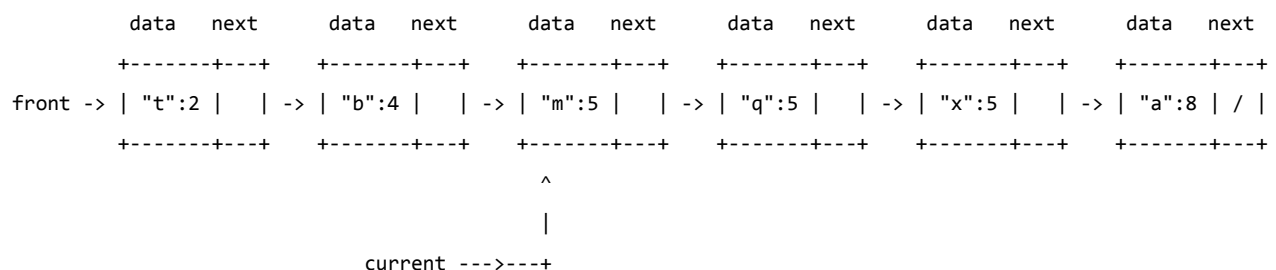> **"t"**, then **"b"**, then **"m"**, then **"q"**, then **"x"**, then **"a"**

You could think of a priority queue as a sorted queue where the elements are sorted by priority, breaking ties by comparing the string elements themselves. But internally the priority queue might not actually store its elements in sorted order; all that matters is that when they are dequeued, they come out in sorted order by priority. An actual priority queue implementation is not required to store its internal data in any particular order, so long as it dequeues its elements in increasing order of priority. As we will see, this difference between the external expected behavior of the priority queue and its true internal state can lead to interesting differences in implementation.
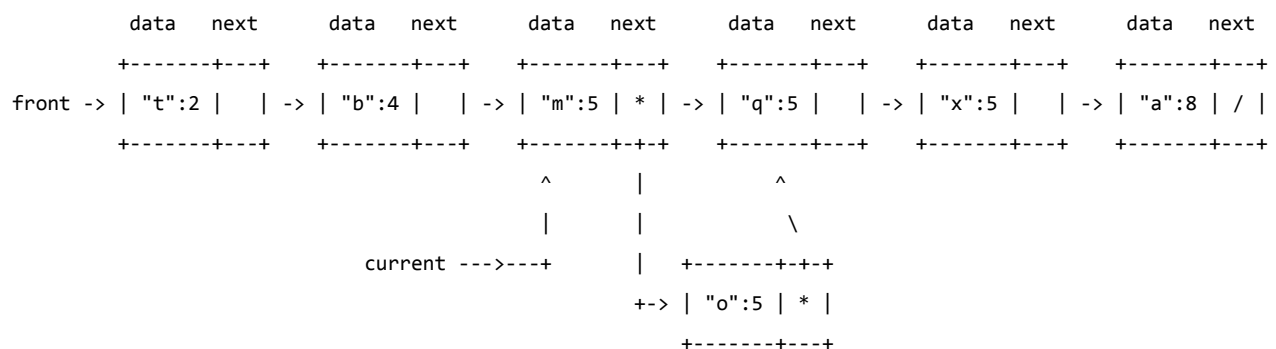
PRIORITY QUEUE IMPLEMENTATIONS

1) SLinkedPriorityQueue: The first priority queue implementation you will write uses a **sorted linked list** as its internal data storage. This class is only allowed to have a single private member variable inside it: a pointer to the front of your list. The elements of the linked list are stored in **sorted** order internally. As new elements are enqueued, you should add them at the appropriate place in the linked list so as to maintain the sorted order. When dequeuing, you do not need to search the linked list to find the smallest element and remove/return it; it is always at the front of the list. The following is a diagram of the internal linked list state of a SLinkedPriorityQueue after enqueuing the elements listed on the previous page:

```
            data    next        data    next        data    next        data    next        data    next        data    next

          +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+
front ->  | "t":2 |   |  -> | "b":4 |   |  -> | "m":5 |   |  -> | "q":5 |   |  -> | "x":5 |   |  -> | "a":8 | / |
          +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+
```

The hardest part of this implementation is inserting a new node in the proper place when enqueue is called. You must look for the proper insertion point by finding the first element whose priority is at least as large as the new value to insert, breaking ties by comparing the strings. Remember that, as shown in class, you must often stop one node early so that you can adjust the next pointer of the preceding node. For example, if you were going to insert the value "o" with priority 5 into the list shown above , your code should iterate until you have a pointer to the node containing "m", as shown below:

```
            data    next        data    next        data    next        data    next        data    next        data    next

          +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+
front ->  | "t":2 |   |  -> | "b":4 |   |  -> | "m":5 |   |  -> | "q":5 |   |  -> | "x":5 |   |  -> | "a":8 | / |
          +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+
                                                      ^
                                                      |
                                  current --->---+
```

Once the current pointer shown above points to the right location, you can insert the new node as shown below:

```
            data    next        data    next        data    next        data    next        data    next        data    next

          +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+       +-------+---+
front ->  | "t":2 |   |  -> | "b":4 |   |  -> | "m":5 | * |  -> | "q":5 |   |  -> | "x":5 |   |  -> | "a":8 | / |
          +-------+---+       +-------+---+       +-------+-+-+       +-------+---+       +-------+---+       +-------+---+
                                                      ^     |           ^
                                                      |     |            \
                                  current --->---+     |   +-------+-+-+
                                                       +-> | "o":5 | * |
                                                           +-------+---+
```

We supply you with a ListNode structure that is a small object representing a single node of the linked list. Each ListNode stores a string value and integer priority in it, and a pointer to a next

node. You should use this structure to store the elements of your priority queue along with their priorities.

Your list is **not allowed** to store an integer size member variable; you must use the presence of a NULL next pointer to figure out where the end of the list is and how long it is.

2) `USLinkedPriorityQueue`**:** The second priority queue implementation you will write uses an **unsorted linked list** as its internal data storage. This class is only allowed to have a single private member variable inside it: a pointer to the front of your list. The elements of the linked list are stored in **unsorted** order internally. As new elements are enqueued, you should *add at the front.* When dequeuing, you need to *search the linked list* to find the smallest element and remove/return it; it could be anywhere. The following is a diagram of the internal linked list state of a `USLinkedPriorityQueue` after enqueuing the elements listed on the previous page:

```
          data   next      data   next      data   next      data   next      data   next      data   next
        +-------+---+     +-------+---+     +-------+---+     +-------+---+     +-------+---+     +-------+---+
front ->| "x":5 |   | -> | "b":4 |   | -> | "a":8 |   | -> | "m":5 |   | -> | "q":5 |   | -> | "t":2 | / |
        +-------+---+     +-------+---+     +-------+---+     +-------+---+     +-------+---+     +-------+---+
```

We supply you with a `ListNode` structure that is a small object representing a single node of the linked list. Each `ListNode` stores a string value and integer priority in it, and a pointer to a next node. You should use this structure to store the elements of your priority queue along with their priorities.

Your list is **not allowed** to store an integer size member variable; you must use the presence of a NULL next pointer to figure out where the end of the list is and how long it is.

3) `HeapPriorityQueue`: The third priority queue implementation you will write uses a special array structure called a binary heap as its internal data storage. The only private member variables this class is allowed to have inside it are a pointer to your internal array of elements, and integers for the array's capacity and the priority queue's size.

As discussed in lecture, a binary heap is an unfilled array that maintains a "heap ordering" property where each index i is thought of as having two "child" indexes, i * 2 and i * 2 + 1, and where the elements must be arranged such that "parent" indexes always store more urgent priorities than their "child" indexes. To simplify the index math, we will leave index 0 blank and start the data at an overall parent "root" or "start" index of 1. One very desirable property of a binary heap is that the most urgent-priority element (the one that should be returned from a call to peek or dequeue) is always at the start of the data in index 1. For example, the six elements listed in the previous pages could be put into a binary heap as follows. Notice that the most urgent element, "t":2, is stored at the root index of 1.

```
index       0       1       2       3       4       5       6       7       8       9
          +-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+
value     |       | "t":2 | "m":5 | "b":4 | "x":5 | "q":5 | "a":8 |       |       |       |
          +-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+
```

```
size    = 6
capacity = 10
```

As discussed in lecture, adding (enqueuing) a new element into a heap involves placing it into the first empty index (7, in this case) and then "bubbling up" or "percolating up" by swapping it with its parent index (i/2) so long as it has a more urgent (lower) priority than its parent. We use integer division, so the parent of index 7 = 7/2 = 3. For example, if we added "y" with priority 3, it would first be placed into index 7, then swapped with "b":4 from index 3 because its priority of 3 is less than b's priority of 4. It would not swap any further because its new parent, "t":2 in index 1, has a lower priority than y. So the final heap array contents after adding "y":3 would be:

```
index       0       1       2       3       4       5       6       7       8       9
        +-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+
value   |       | "t":2 | "m":5 | "y":3 | "x":5 | "q":5 | "a":8 | "b":4 |       |       |
        +-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+
size    = 7
capacity = 10
```

Removing (dequeuing) the most urgent element from a heap involves moving the element from the last occupied index (7, in this case) all the way up to the "root" or "start" index of 1, replacing the root that was there before; and then "bubbling down" or "percolating down" by swapping it with its more urgent-priority child index (i*2 or i*2+1) so long as it has a less urgent (higher) priority than its child. For example, if we removed "t":2, we would first swap up the element "b":4 from index 7 to index 1, then bubble it down by swapping it with its more urgent child, "y":3 because the child's priority of 3 is less than b's priority of 4. It would not swap any further because its new only child, "a":8 in index 6, has a higher priority than b. So the final heap array contents after removing "t":2 would be:

```
index       0       1       2       3       4       5       6       7       8       9
        +-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+
value   |       | "y":3 | "m":5 | "b":4 | "x":5 | "q":5 | "a":8 |       |       |       |
        +-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+
size    = 6
capacity = 10
```

A key benefit of using a binary heap to represent a priority queue is efficiency. The common operations of enqueue and dequeue take only O(log N) time to perform, since the "bubbling" jumps by powers of 2 every time. The peek operation takes only O(1) time since the most urgent-priority element's location is always at index 1.

If nodes ever have a tie in priority, break ties by comparing the strings themselves, treating strings that come earlier in the alphabet as being more urgent (e.g. **"a"** before **"b"**). Compare strings using the standard relational operators like <, <=, >, >=, ==, and !=. Do not make assumptions about the lengths of the strings.

Changing the priority of an existing value involves looping over the heap to find that value, then once you find it, setting its new priority and "bubbling up" that value from its present location, somewhat like an enqueue operation.

For heap PQs, when the array becomes full and has no more available indexes to store data, you must resize it to a larger array. **Your larger array should be a multiple of the old array size**, **such as double the size**. **You must not leak memory**; free all dynamically allocated arrays created by your class.

PRIORITY QUEUE OPERATIONS
Each of your three priority queue implementations must support all of the following operations.

Member	Description
pq.enqueue(*value, priority*);	In this function you should add the given string value into your priority queue with the given priority. Duplicates are allowed. Any string is a legal value, and any integer is a legal priority; there are no invalid values that can be passed.
pq.dequeue()	In this function you should remove the element with the most urgent priority from your priority queue, and you should also return it. You should throw a string exception if the queue does not contain any elements.
pq.peek()	In this function you should return the string element with the most urgent priority from your priority queue, without removing it or altering the state of the queue. You should throw a string exception if the queue does not contain any elements.
pq.peekPriority()	In this function you should return the integer priority that is most urgent from your priority queue (the priority associated with the string that would be returned by a call to peek), without removing it or altering the state of the queue. You should throw a string exception if the queue does not contain any elements.
pq.changePriority(*value, newPriority*);	In this function you will modify the priority of a given existing value in the queue. The intent is to change the value's priority to be more urgent (smaller integer) than its current value. If the given value is present in the queue and already has a more urgent priority to the given new priority, or if the given value is *not* already in the queue, your function should throw a string exception. If the given value occurs multiple times in the priority queue, you should alter the priority of the first occurrence you find when searching your internal data from the start.
pq.isEmpty()	In this function you should return true if your priority queue does not contain any elements and false if it does contain at least one element.
pq.size()	In this function you should return the number of elements in your priority queue.

pq.clear();	In this function you should remove all elements from the priority queue.
out << *pq*	You should write a << operator for printing your priority queue to the console. The elements can print out <u>in any order</u> and must be in the form of *"value"*:*priority* with {} braces, such as {"t":2, "b":4, "m":5, "q":5, "x":5, "a":8} . The PQEntry and ListNode structures both have << operators that may be useful. Your formatting and spacing should match *exactly*. Do not place a \n or endl at the end.

The headers of every operation must match those specified above. **Do not change** the parameters or function names.

Constructor/destructor: Each class must also define a parameterless constructor. If the implementation allocates any dynamic memory, you must ensure that there are no memory leaks by freeing any allocated memory at the appropriate time. This will mean that you will need a destructor for classes that dynamically allocate memory.

Helper functions: The members listed on the previous page represent a large fraction of each class's behavior. But you should add other members to help you implement all of the appropriate behavior. Any other member functions you provide must be private. Remember that each member function of your class should have a clear, coherent purpose. You should provide private helper members for common repeated operations. Make a member function and/or parameter const if it does not perform modification of the object's state.

Member variables: We have already specified what member variables you should have. Here are some other constraints:
- Don't make something a member variable if it is only needed by one function. Make it local. Making a variable into a data member that could have been a local variable or parameter will hurt your Style grade.
- All data member variables inside each of your classes should be private.

OTHER IMPLEMENTATION DETAILS
Here are a few other constraints we expect you to follow that don't fit neatly into any other section.
- The HeapPriorityQueue has some similar aspects to the ArrayList written in lecture. It's completely fine to borrow any of the code from that lecture and use it in your Array PQ as needed. *(Cite it as a source.)*
- The US/SLinkedPriorityQueue's operations should not make unnecessary passes over the linked list. For example, when enqueuing an element in sorted order, a poor implementation would be to traverse the entire list once to count its size and to find the proper spot to insert, and then make a second traversal to get back to the spot to insert and add the new element. Do not make such multiple passes. Also, keep in mind that your US/SLinkedPriorityQueue is **<u>not</u> allowed** to store an integer size member

variable; you must use the presence of a NULL next pointer to figure out where the end of the list is and how long it is.

- The `HeapPriorityQueue` must implement its operations efficiently using the "bubbling" or "percolating" described in this handout. It is important that these operations run in O(log N) time.

- Duplicates are allowed in your priority queue, so be mindful of this. For example, the changePriority operation should affect only a single occurrence of a value (the first one found). If there are other occurrences of that same value in the queue, a single call to changePriority shouldn't affect them all.

- You are not allowed to use a sort function to arrange the elements of any collection, nor are you allowed to create any temporary or auxiliary data structures inside any of your priority queue implementations. They must implement all of their behavior using only their primary internal data structure as specified.

- You will need pointers for several of your implementations, but you should not use pointers-to-pointers (for example, `ListNode**`) or references to pointers (e.g. `ListNode*&`).

- You should not create any more `ListNode` objects than necessary. For example, if a US/SLinkedPriorityQueue contains 6 elements, there should be exactly 6 `ListNode` objects in the chain, no more, no less. You shouldn't, say, have a seventh empty "dummy" node at the front or back of the list that serves only as a marker. You can declare as many local variable pointers to `ListNode`s as you like.

DEVELOPMENT TIPS AND STRATEGY

Order of Implementation: While you are free to implement the three priority queues in any order you wish, we recommend that you implement them in the order we specified: Linked Lists, then Heap. This goes from simplest to most difficult.

Draw pictures. Remember the first rule of linked lists: draw pictures of linked lists! Draw pictures of the before, during, and after state each of the operations you perform on it. Manipulating linked lists can be tricky, but if you have a picture in front of you as you're coding it can make your job substantially easier.

Don't panic. You will be doing a lot of pointer gymnastics in the course of this assignment, and you will almost certainly encounter a crash in the course of writing your program. If your program crashes, resist the urge to immediately make changes to your code. Instead, look over your code methodically. Use the debugger to step through the code one piece at a time, or use the provided testing harness to execute specific commands on the priority queue. The bug is waiting there to be found, and with persistence you will find it. If your program crashes with a specific error message, try to figure out exactly what that message means. Don't hesitate to get in touch with your section leader, and feel free to stop by the LaIR or office hours.

Testing: Extensively test your program. There is a provided test harness that will help you with testing. You will need to edit it specifically to perform the performance analysis, but the version provided in the starter code already runs many of the tests you'll want to do while you're still

implementing and debugging. Be sure to take advantage of that resource! Of course, we won't guarantee that our test harness will catch all your bugs.

EMPIRICAL PERFORMANCE ANALYSIS

A major purpose of doing all these different implementations is to compare and contrast how difficult/complicated they are to code (which correlates with potential for bugs), and how they perform. You've no doubt thought about the former as you've worked on the assignment. For this final part of the assignment, you will run experiments to formally investigate the latter.

A test harness is included with the code (pqueue-main.cpp). Edit this as you see fit (no need to perfectly match the original user interaction, etc.) so that it will generate random strings and enqueue them into your priority queues, then dequeue them again, and time the operations. Vary the "N" (number of strings enqueued/dequeued) so you can plot a few points and begin to see the expected big-O behavior for each implementation. First just play around with the N for a while to see where the most relevant or "interesting" range of values to use is (for the purposes of observing Big-O behavior and comparing implementations), and then you should do your official experiment runs.

- You may want to in some way reuse or copy and slightly rewrite the bulk enqueue code that is already in the provided test harness.
- If your N is too small, the timing that you do will be influenced too much by background noise on the system and won't elucidate the actual performance of your data structures.
- If your N is too large, the performance differences you observe may have more to do with whether you fit in cache or main memory than the performance of your data structures. (Though this would be a fascinating thing to observe and document!)
- One way to reduce background noise, which you should definitely plan on doing, is running the test many times to "average out" occasional variations that occur (e.g., when your computer decides to do something other than execute your testing code). For example, if you want to run on N=1000, do something like this pseudocode (this repeats tests 100 times but you might try larger and smaller values to see where you get the best noise filtering without wasting time unnecessarily):

  ```
  Generate 1000 random strings
  Start timer
  For (i = 0 to 99)
      Create empty PriorityQueue
      Enqueue 1000 strings
  End timer
  ```
- To do the timings, use the Stanford Library Timer class:
  http://stanford.edu/~stepp/cppdoc/Timer-class.html

Your report should include:
- 1 plot showing timing results for ENQUEUE (plot all implementations in the same figure for comparison). You should plot at least 5 data points (5 values of N) per implementation, but you could do many, many more if you want.

- 1 plot showing timing results for DEQUEUE (plot all implementations in the same figure for comparison). You should plot at least 5 data points (5 values of N) per implementation, but you could do many, many more if you want.
- A short paragraph explaining how you set up your tests (what is the range of N values you tried, how many test repetitions were needed to effectively reduce noise, etc.).
- A short paragraph drawing conclusions about your results, including remarks on whether you were able to observe the expected Big-O behavior. It's ok to say that you were not able to get a particularly clear picture of this, but you should describe your efforts and speculate as to what the roadblocks might be.

A key word for this part of the assignment is "simple." It is not required to dazzle us with the graphic design of the plots or anything like that.

STYLE GUIDELINES SPECIFIC TO THIS ASSIGNMENT
In general, items mentioned in the "Implementation and Grading" from the previous assignment(s) specs also apply here. Please refer to those documents as needed. Note the instructions in the previous assignments about procedural decomposition, variables, types, parameters, value vs. reference, and commenting. Don't forget to cite any sources you used in your comments. Refer to the course Style Guide for a more thorough discussion of good coding style.

Commenting: Since all of the queue classes have the same public members, we will allow you to comment the public member functions (enqueue, dequeue, etc.) a single time in `SLinkedPriorityQueue`, and then in the other classes you can simply write, "See `SLinkedPriorityQueue.h` for documentation member functions." But we do expect you to put a descriptive comment header on each queue class file and explain that implementation and its pros and cons. Also put a comment atop every member function that states its Big-Oh. For example, in a `HeapPriorityQueue` the peek operation runs in constant time, so you should put a comment on that function that says "O(1)."

Redundancy: Redundancy is another major grading focus; avoid repeated logic as much as possible. Your classes will be graded on whether you make good choices about what members it should have, and other factors such as which are public vs. private, and const-correctness, and so on. You may find that there are some operations that you have to repeat in all of your classes, like checking whether a queue is empty before dequeuing from it. We do **not** require you to reduce redundancy across multiple classes (for example, copy/paste code between the two linked lists rather than attempting to unify them somehow); but we do expect you to remove redundancy within a single class. If one implementation has a common operation, make a private helper function and call it multiple times in that file.

Follow the Honor Code on this assignment. Submit your own (pair's) work; do not look at others' solutions. Cite sources. Do not give out your solution or place it on a public web site or forum. The Stanford C++ library includes a priority queue similar to the `HeapPriorityQueue` you must

implement. We ask that you do not look at that file's source code because it is too similar to what you are asked to do here. You must solve the problem yourself.

EXTRA CREDIT IDEAS

A good extra would be to write a PQ implementation(s) beyond those required by this assignment. Below are some ideas. The ideas that are especially recommended for CS106X students are marked with **. Remember to include your extra data structures in your empirical analysis. Especially for Binomial Heap, you should see the dazzling performance results of your efforts!

- **Map of Queues:** Use a Map with integer priorities as its keys and queues of strings as the values associated with those keys. This puts all elements with the same priority into an inner queue together.
- **Binary Search Tree:** We talked about BST as an implementation of Map, not PriorityQueue, but it is possible to use BST for PriorityQueue.
- ****Doubly-Linked List:** Rewrite one or both of your linked list implementations as a doubly-linked list (a list whose nodes have both prev and next pointers). This makes it easy to move both forward and backward in the list, which can speed up certain operations like changing priority. Also declare both a front and a back pointer, which point to the first and last element in the list, respectively. Implementing a doubly-linked list is a classic "bucket list" thing to do in learning how to program.
- ****Binomial Heap:** A binomial heap is a variation on the binary heap already described that is especially efficient for certain common priority queue operations. See separate handout on class web site describing this idea. Binomial Heap implementation is the experience of legend in CS106X. Super fun data structure!! (for some definition of fun ☺)

Another good idea is to add operations to each heap beyond those specified:

- Merge: Write a member function that accepts another priority queue of the same type and adds all of its elements into the current priority queue. Do this merging "in place" as much as possible; for example, if you are merging two linked list PQs, directly connect the node pointers of one to the other as appropriate.
- Deep Copy: Make your priority queues properly support the = assignment statement, copy constructor, and deep copying. See the C++ "Rule of Three" and follow that guideline in your implementation.
- Iterator: Write a class that implements the STL iterator and a begin and end function in your priority queues, which would enable "for-each" over your PQ. This requires knowledge of the C++ STL library.
- Other: Do you have an interesting idea for an extra feature? Ask the head TA or instructor.

Submitting with extra features: If you complete any extras, please list them in your comment headings. Also please submit your program twice: first without extra features (or with them disabled), and a second time with the extensions.