

Lecture Notes on Queues

15-122: Principles of Imperative Computation
Frank Pfenning

Lecture 9
February 8, 2011

1 Introduction

In this lecture we introduce *queues* as a data structure and *linked lists* that underly their implementation. In order to implement them we need *recursive types*, which are quite common in the implementation of data structures.

2 Linked Lists

Linked lists are a common alternative to arrays in the implementation of data structures. Each item in a linked list contains a data element of some type and a *pointer* to the next item in the list. It is easy to insert and delete elements in a linked list, which is not a natural operation on arrays. On the other hand access to an element in the middle of the list is usually $O(n)$, where n is the length of the list.

An item in a linked list consists of a struct containing the data element and a pointer to another linked list. This gives rise to the following definition:

```
struct list {
    string data;
    struct list* next;
};
typedef struct list* list;
```

This definition is an example of a *recursive type*. A struct of this type contains a pointer to another struct of the same type, and so on. We usually use the special element of type `t*`, namely `NULL`, to indicate that we have reached the end of the list. Sometimes (as will be the case for queues introduced next), we can avoid the explicit use of `NULL` and obtain more elegant code. The type definition is there to create the type name `list`, which stands for a pointer to a `struct list`.

There are some restriction on recursive types. For example, a declaration such as

```
struct infinite {
    int x;
    struct infinite next;
}
```

would be rejected by the C0 compiler because it would require an infinite amount of space. The general rule is that a struct can be recursive, but the recursion must occur beneath a pointer or array type, whose values are addresses. This allows a finite representation for values of the struct type.

We don't introduce any general operations on lists; let's wait and see what we need where they are used. Linked lists as we use them here are a *concrete type* which means we do *not* construct an interface and a layer of abstraction around them. When we use them, we know about and exploit their precise internal structure. This is contrast to *abstract types* such as queues or stacks (see next lecture) whose implementation is hidden behind an interface, exporting only certain operations. This limits what clients can do, but it allows the author of a library to improve its implementation without having to worry about breaking client code. Concrete types are cast into concrete once and for all.

3 The Queue Interface

A *queue* is a data structure where we add elements at the back and remove elements from the front. In that way a queue is like "waiting in line": the first one to be added to the queue will be the first one to be removed from the queue. This is also called a FIFO (First In First Out) data structure. Queues are common in many applications. For example, when we read a book from a file as in Assignment 2, it would be natural to store the the words in a queue so that when we are finished reading the file the words are in the order they appear in the book. Another common example are

buffers for network communication that temporarily store packets of data arriving on a network port. Generally speaking, we want to process them in the order that they arrive.

Before we consider the implementation to a data structure it is helpful to consider the interface. We then program against the specified interface. Based on the description above, we require the following functions:

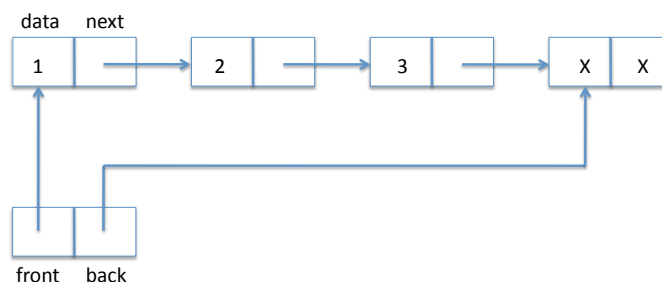
```
typedef struct queue* queue;

bool queue_empty(queue Q);      /* 0(1), check if queue is empty */
queue queue_new();              /* 0(1), create new empty queue */
void enq(queue Q, string s);    /* 0(1), add item at back */
string deq(queue Q);           /* 0(1), remove item from front */
```

We can write out this interface without committing to an implementation of queues. After the type definition we know only that a queue will be implemented as a pointer to a struct `queue`.

4 The Queue Implementation

Here is a picture of the queue data structure the way we envision implementing it, where we have elements 1, 2, and 3 in the queue.



A queue is implemented as a struct with a `front` and `back` field. The `front` field points to the front of the queue, the `back` field points to the back of the queue. In arrays, we often work with the length which is one greater than the index of the last element in the array. In queues, we use a similar strategy, making sure the back pointer points to one element past the end of the queue. Unlike arrays, there must be something in memory for the pointer to refer to, so there is always one extra element at the end of the

queue which does not have valid data or next pointer. We have indicated this in the diagram by writing X.

The above picture yields the following definition.

```
struct queue {
    list front;
    list back;
};
```

When does a struct of this type represent a valid queue? In fact, whenever we define a new data type representation we should first think about the data structure invariants. Making these explicit is important as we think about and write the pre- and postconditions for functions that implement the interface.

What we need here is if we follow `front` and then move down the linked list we eventually arrive at `back`. We call this a *list segment*. We also want both `front` and `back` not to be `NULL` so it conforms to the picture, with one element already allocated even if the queue is empty.

```
bool is_queue(queue Q) {
    if (Q == NULL) return false;
    if (Q->front == NULL || Q->back == NULL) return false;
    return is_segment(Q->front, Q->back);
}
```

Next, the code for checking whether two pointers delineate a list segment. When both `start` and `end` are `NULL`, we consider it a valid list segment, even though this will never come up for queues. It is a common code pattern for working with linked lists and similar data representation to have a pointer variable, here called *p*, that is updated to the next item in the list on each iteration until we hit the end of the list.

```
bool is_segment(list start, list end) {
    list p = start;
    while (p != end) {
        if (p == NULL) return false;
        p = p->next;
    }
    return true;
}
```

Here we stop in two situations: if $p = \text{null}$, then we cannot come up against *end* any more because we have reached the end of the list and we return

false. The other situation is if we find *end*, in which case we return true since we have a valid list segment. This function may not terminate if the list contains a cycle. We will address this issue in the next lecture; for now we assume all lists are acycle.

To check if the queue is empty, we just compare its front and back. If they are equal, the queue is empty; otherwise it is not. We require that we are being passed a valid queue. Generally, when working with a data structure, we should always require and ensure that its invariants are satisfied in the pre- and post-conditions of the functions that manipulate it.

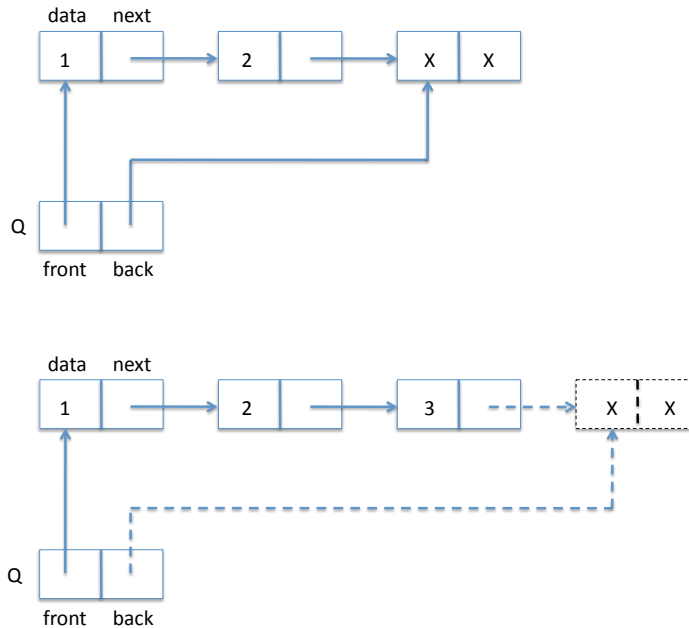
```
bool queue_empty(queue Q)
//@requires is_queue(Q);
{
    return Q->front == Q->back;
}
```

To obtain a new empty queue, we just allocate a list struct and point both front and back of the new queue to this struct. We do not initialize the list element because its contents are irrelevant, according to our representation.

```
queue queue_new()
//@ensures is_queue(\result);
//@ensures queue_empty(\result);
{
    queue Q = alloc(struct queue);
    list l = alloc(struct list);
    Q->front = l;
    Q->back = l;
    return Q;
}
```

To enqueue something, that is, add a new item to the back of the queue, we just write the data (here: a string) into the extra element at the back, create a new back element, and make sure the pointers updated correctly. You should draw yourself a diagram before you write this kind of code. Here is a before-and-after diagram for inserting "3" into a list. The new or

updated items are dashed in the second diagram.



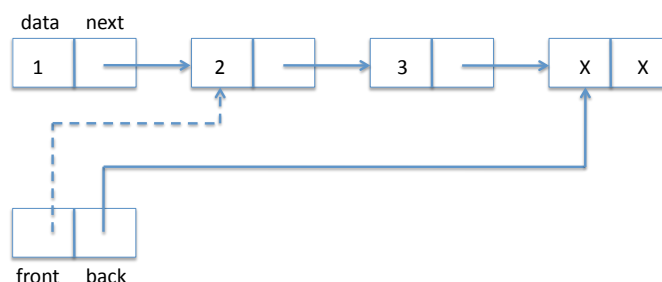
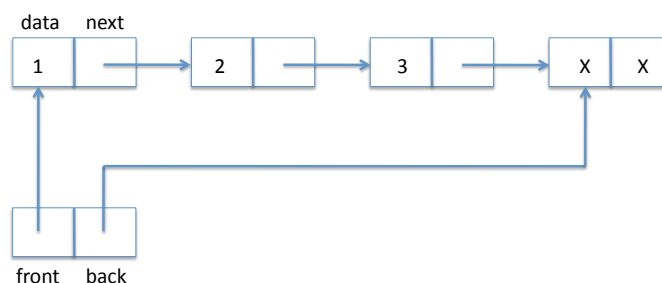
Another important point to keep in mind as you are writing code that manipulates pointers is to make sure you perform the operations in the right order, if necessary saving information in temporary variables.

```
void enq(queue Q, string s)
//@requires is_queue(Q);
//@ensures is_queue(Q);
{
    list l = alloc(struct list);
    Q->back->data = s;
    Q->back->next = l;
    Q->back = l;
}
```

The invariant marked *internal* is one that only makes sense in the context of the present implementation, but would not be meaningful to a client.

Finally, we have the dequeue operation. For that, we only need to change the front pointer, but first we have to save the dequeued element

in a temporary variable so we can return it later. In diagrams:



And in code:

```
string deq(queue Q)
//@requires is_queue(Q);
//@requires !queue_empty(Q);
//@ensures is_queue(Q);
{
    assert(!queue_empty(Q));
    string s = Q->front->data;
    Q->front = Q->front->next;
    return s;
}
```

We included an explicit `assert` statement with an error message so that if `deq` is called with an empty queue we can issue an appropriate error message. Unlike the precondition of the function, this will *always* be checked, which is good practice when writing library code that might be called incorrectly from the outside.

We do not always check whether the given queue is valid for two reasons. First, it takes $O(n)$ time when there are n elements in the queue, so

dequeuing and enqueueing would no longer be constant time. Second, if the client respects the interface and manipulates the data structure only through the given interface, then it should not be possible to construct an invalid queue. On the other hand, it is perfectly possible to construct an empty queue and mistakenly hand it to the `deq` function, so we check this condition explicitly.

An interesting point about the dequeue operation is that we do not explicitly deallocate the first element. If the interface is respected there cannot be another pointer to the item at the front of the queue, so it becomes *unreachable*: no operation of the remainder of the running programming could ever refer to it. This means that the garbage collector of the C0 runtime system will recycle this list item when it runs short of space.