

## Node-Based Structures More on Linked Lists

---

CS 311 Data Structures and Algorithms

Lecture Slides

Monday, November 2, 2009

Glenn G. Chappell

Department of Computer Science

University of Alaska Fairbanks

**CHAPPELLG@member.ams.org**

© 2005–2009 Glenn G. Chappell

# Unit Overview

## Handling Data & Sequences

---

### Major Topics

- ✓ ■ Data abstraction
- ✓ ■ Introduction to Sequences
- ✓ ■ Smart arrays
  - ✓ ■ Array interface
  - ✓ ■ Basic array implementation
  - ✓ ■ Exception safety
  - ✓ ■ Allocation & efficiency
  - ✓ ■ Generic containers
- Linked Lists
  - Node-based structures
  - More on Linked Lists
- Sequences in the C++ STL
- Stacks
- Queues

## Review

### Exception Safety [1/2]

---

**Safety:** Does function ever signal an error condition, and if so:

- Are data left in a usable state?
- Do we know something about that state?
- Are resource leaks avoided?

#### **Basic Guarantee**



Minimum  
standard

- Data remain in a usable state, and resources are never leaked, even in the presence of exceptions.

Each guarantee  
includes the  
previous one.

#### **Strong Guarantee**



Preferred

- If the operation throws an exception, then it makes no changes that are visible to the client code.

#### **No-Throw Guarantee**



Required in some  
special situations

- The operation never throws an exception.

## Review

### Exception Safety [2/2]

---

To make sure code is exception-safe:

- Look at **every** place an exception might be thrown.
- For each, make sure that, if an exception is thrown, either
  - we terminate normally and meet our postconditions, or
  - we throw and meet our guarantees.

A bad design can force us to be unsafe.

- Thus, good design is part of exception safety.
- An often helpful idea is that every module has exactly one well defined responsibility (the **Single Responsibility Principle**).
- In particular: A non-const member function should not return an object by value.

## Review

### Allocation & Efficiency [1/2]

---

An operation is **amortized constant time** if  $k$  operations require  $O(k)$  time.

- Thus, over *many consecutive operations*, the operation averages constant time.
- Not the same as constant-time average case (which averages over *all possible inputs*)
- Quintessential amortized-constant-time operation: insert-at-end for a well written smart array.
- Amortized constant time is not something we can easily compare with (say) logarithmic time.

## Review

### Allocation & Efficiency [2/2]

---

#### Improving `SmArray`

- Our original design did not allow for efficient insert-at-end.
  - Reallocate-and-copy would happen every time.
- The revised design had three data members: size, capacity, and the array pointer.
- Having a “capacity” member allows us to keep a record of how much memory is allocated. Then we can allocate extra so that we do not need to reallocate every time.
- Result: We can do amortized-constant-time insert-at-end.

## Review

### Generic Containers [1/2]

---

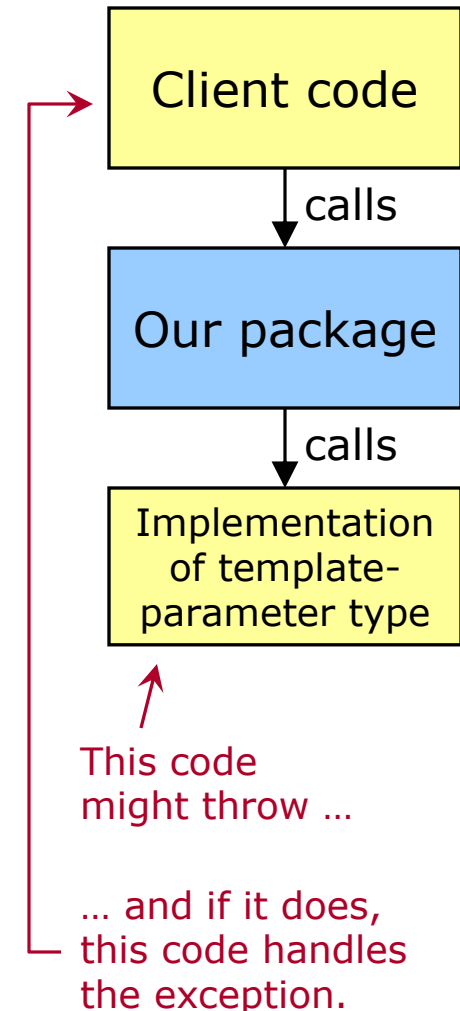
A **generic container** is a container that can hold a client-specified data type.

- In C++ we usually implement a generic container using a **class template**.

A function that allows exceptions thrown by a client's code to propagate unchanged, is said to be **exception-neutral**.

When exception-neutral code calls a client-provided function that may throw, it does one of two things:

- Call the function outside a try block, so that any exceptions terminate our code immediately.
- Or, call the function inside a try block, then catch all exceptions, do any necessary clean-up, and re-throw.



## Review

### Generic Containers [2/2]

---

We can use catch-all, clean-up, re-throw to get both exception safety and exception neutrality.

```
arr = new MyType[size];  
try  
{  
    std::copy(a, a+size, arr);  
}  
catch (...)  
{  
    delete [] arr;  
    throw;  
}
```

Called outside any `try` block. If this fails, we exit immediately, throwing an exception.

Called inside a `try` block. If this fails, we need to deallocate the array before exiting.

This helps us meet the Basic Guarantee (also the Strong Guarantee if this function does nothing else).

This makes our code exception-neutral.

# Node-Based Structures

## Introduction [1/2]

---


Our fundamental building block for data structures has been the **array**.



- Items are stored in contiguous memory locations.
- Look-up operations are usually very fast.
- Operations that require rearrangement (insert, remove, etc.) can be slow.

Now we begin looking at data structures built out of **nodes**.

- A *node* is generally a small block of memory that is referenced via a pointer, and which may reference other nodes via pointers.




- Node-based structures do not necessarily store data in contiguous memory.
  - Memory-management changes significantly.
- To find a node, we follow a chain of pointers. Look-up can be slow.
- Operations that require rearrangement can be very fast.

## Node-Based Structures Introduction [2/2]

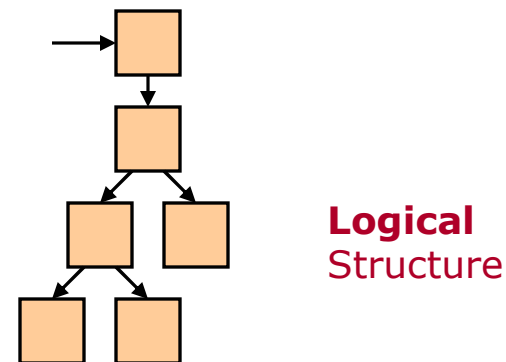
---

When we draw pictures of node-based data structures, the positions of nodes in the picture usually have nothing to do with their positions in memory.

For example, if a structure is stored like this ...




... then we might draw it like this:



## Node-Based Structures Implementation — Storage Options



---

We normally store nodes in separately allocated blocks of memory.



However, we might allocate an **array of nodes**.

- Replace pointers with array indices, if desired.



Both of these have the above **logical** structure.

This is still a node-based data structure.

- It still has most of the pros & cons of node-based structures.
- The main differences involve **memory management**: who does it & when it gets done.

## Node-Based Structures Implementation — Classes

---

When we define a class to implement a node-based data structure, we may wish to write several classes:

- The main **container class**, representing the structure as a whole.
- A class representing a **node**.
  - This might be a private member of the main class.
  - Typically client code does not deal with this class.
  - If there are different kinds of nodes, then we might want several node classes. Sometimes inheritance can be helpful here.
- Possibly an **iterator** class.
  - Iterators to node-based structures are almost never pointers, because `operator++`, for example, needs to go to the next **logical** node, not the next **physical** memory location.

Despite the multiple classes being defined, we are implementing only a single package.

- Thus, multiple header & source files are generally not necessary.
- We can make all of these classes friends without breaking encapsulation.

## Node-Based Structures Implementation — Pointers & Ownership

---

Think of nodes as resources to be owned & managed.

- Who owns them?
  - Always document ownership (here: in the class invariants).
- Internal pointers in a node-based structure will usually be owning pointers.
  - A node is typically owned by the node that points to it.
  - Thus, a node's destructor should free all nodes that it points to.

This can make destroying a node-based structure **easy**.

- Each node is responsible for destroying the nodes it owns.
- Thus, to destroy the whole structure, all we need to do is destroy the nodes that are not owned by other nodes.
- And there is usually just one of these.

## More on Linked Lists Refresher [1/2]

---


Earlier in the semester, we looked briefly at the simplest node-based structure: a **Linked List**.

- Like an array, a Linked List is a structure for storing a sequence of items.

Array 

3	1	5	3	5	2
---	---	---	---	---	---

- A Linked List is composed of **nodes**. Each has a single data item and a pointer to the next node.



- These pointers are the **only** way to find the next data item. Thus, unlike an array, we cannot quickly skip to (say) the 100th item in a Linked List. Nor can we quickly find the previous item.
- A Linked List is a one-way sequential-access data structure. Thus, its natural iterator is a **forward iterator**, which has only the ++ operator.

## More on Linked Lists Refresher [2/2]


---

### Why not always use (smart) arrays?

- One important reason: we can often insert and remove much faster with a Linked List.


### Inserting

- Inserting an item at a given position in an array is slow-ish.
- Inserting an item at a given position (think “iterator”) in a Linked List is very fast.
- Example: insert a “7” after the bold node.



### Removing

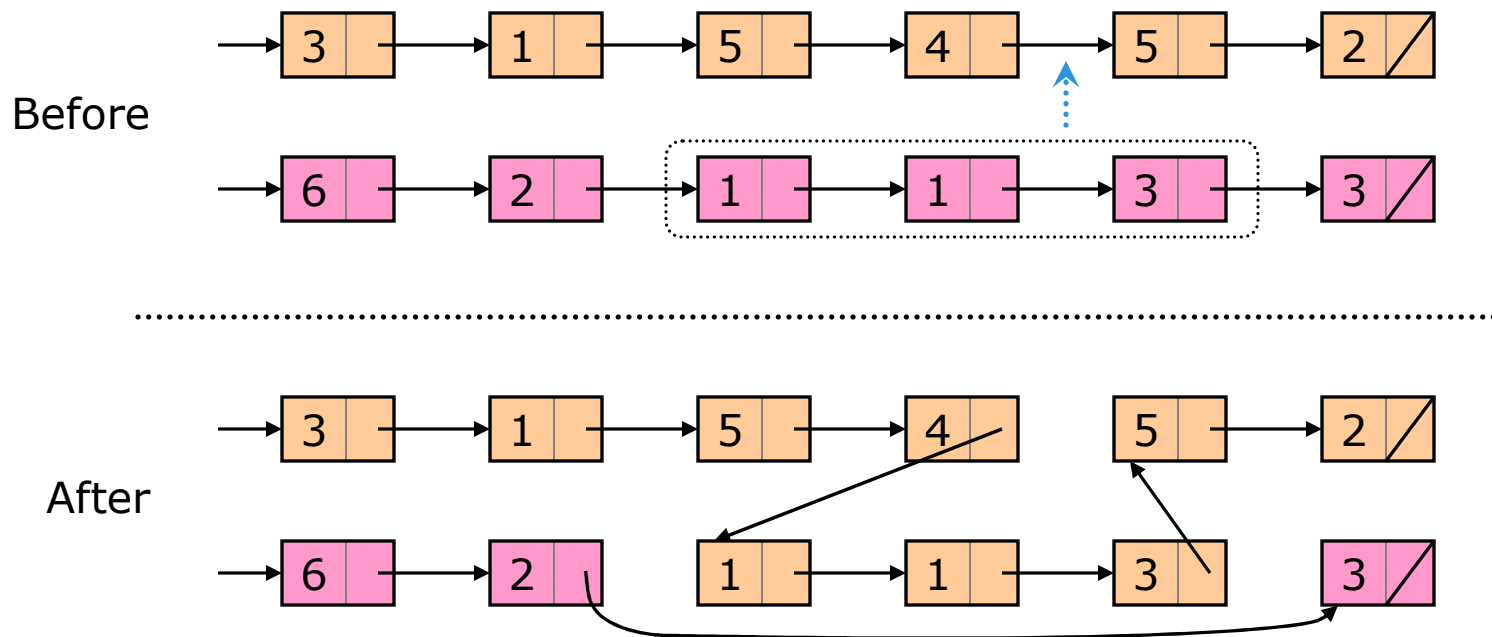
- Removing the item at a given position from an array *is also slow-ish*.
- Removing the item at a given position from a Linked List is very fast.
  - We need an iterator to the *previous* item.
- Example: Remove the item in the bold node.



## More on Linked Lists More Advantages [1/2]

---

With Linked Lists, we can do a fast **splice**:

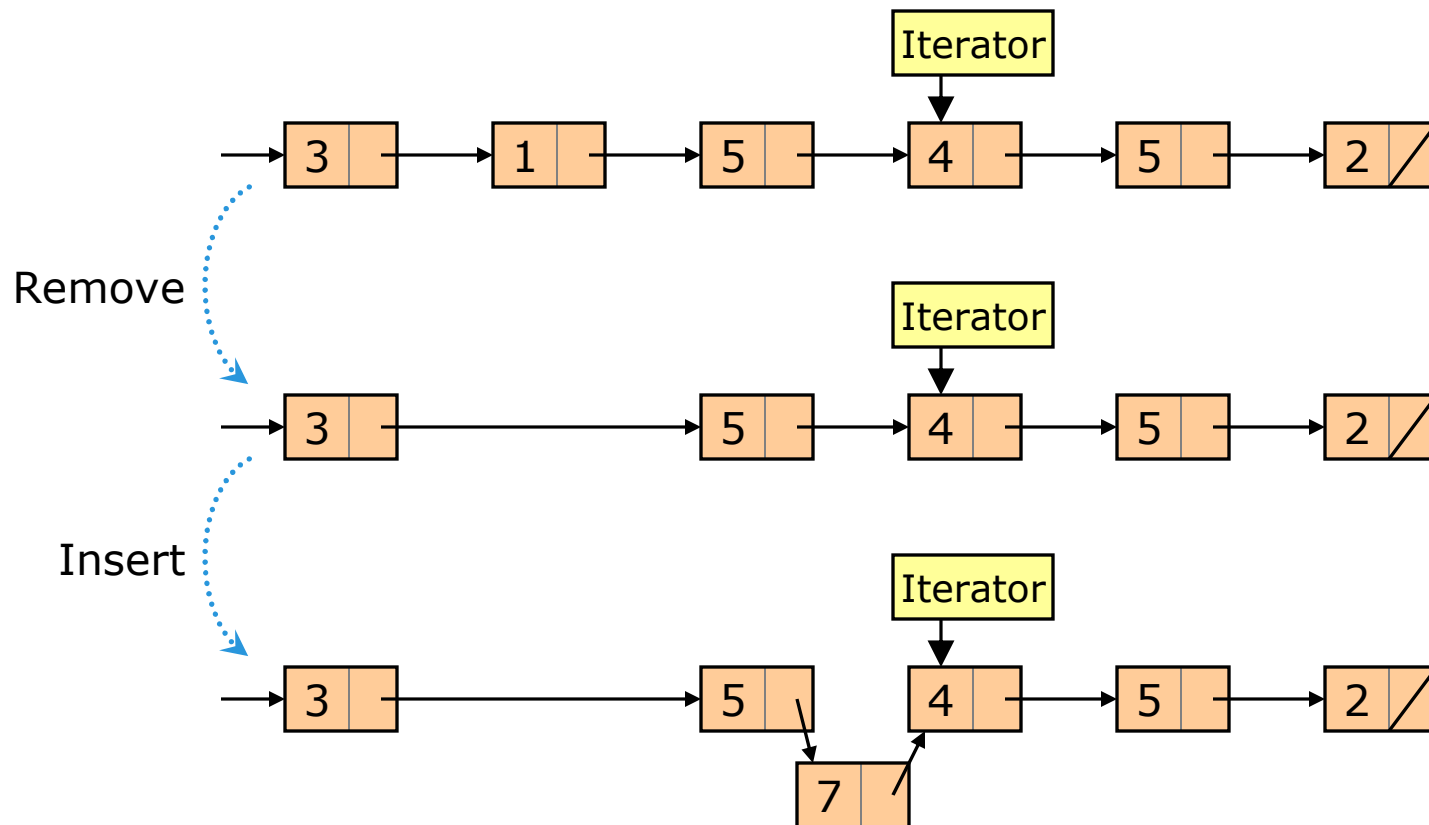


Note: If we allow for efficient splicing, then we cannot efficiently keep track of a Linked List's size.

## More on Linked Lists More Advantages [2/2]

---

With Linked Lists, iterators, pointers, and references to items will always stay valid and never change what they refer to, as long as the Linked List exists — unless we remove or change the item ourselves.



## More on Linked Lists

### Comparison With Arrays [1/4]

---

1. What is the order of each of the following when using (a) a smart-array implementation of a Sequence, and (b) a Linked-List implementation? Assume good design and good algorithms.
  - Look-up an item by index.
  - Search in a sorted Sequence.
  - Search in an unsorted Sequence.
  - Sort a Sequence.
  - Insert an item at a given position.
  - Remove an item at a given position.
  - Splice part of one Sequence into another.
  - Insert item at beginning of Sequence.
  - Remove item at beginning of Sequence.
  - Insert item at end of Sequence.
  - Remove item at end of Sequence.
  - Traverse a Sequence (iterate through all items, in order).
  - Make a copy of an entire Sequence.
  - Swap two Sequences.
2. What other issues arise when comparing the two data structures?

# More on Linked Lists

## Comparison With Arrays [2/4]

	Smart Array	Linked List
<b>Look-up by index</b>	<b><math>O(1)</math></b>	<b><math>O(n)</math></b>
<b>Search sorted</b>	<b><math>O(\log n)</math></b>	<b><math>O(n)</math></b>
Search unsorted	$O(n)$	$O(n)$
Sort	$O(n \log n)$	$O(n \log n)$
<b>Insert @ given pos</b>	<b><math>O(n)</math></b>	<b><math>O(1)^*</math></b>
<b>Remove @ given pos</b>	<b><math>O(n)</math></b>	<b><math>O(1)^*</math></b>
<b>Splice</b>	<b><math>O(n)</math></b>	<b><math>O(1)</math></b>
<b>Insert @ beginning</b>	<b><math>O(n)</math></b>	<b><math>O(1)</math></b>
<b>Remove @ beginning</b>	<b><math>O(n)</math></b>	<b><math>O(1)</math></b>
Insert @ end	$O(1)$ or $O(n)^{**}$ amortized const	$O(1)$ or $O(n)^{***}$
Remove @ end	$O(1)$	$O(1)$ or $O(n)^{***}$
Traverse	$O(n)$	$O(n)$
Copy	$O(n)$	$O(n)$
Swap	$O(1)$	$O(1)$

\*For Singly Linked Lists, we mean inserting or removing just *after* the given position.

- Doubly Linked Lists can help.

\*\* $O(n)$  if reallocation occurs. Otherwise,  $O(1)$ . Amortized constant time.

- Pre-allocation can help.

\*\*\*For  $O(1)$ , need a pointer to the end of the list. Otherwise,  $O(n)$ .

- This is tricky.
- And, for remove @ end, it is basically impossible.
- Doubly Linked Lists can help.

**Find** faster  
with an array

**Rearrange** faster  
with a Linked List

## More on Linked Lists Comparison With Arrays [3/4]

---

### Other Issues

- ☹ Linked Lists use **more memory**.
- ☹ When order is the same, Linked Lists are almost always **slower**.
  - Arrays might be 2–10 times faster.
- ☹ Arrays keep consecutive items in **nearby memory locations**.
  - Many algorithms have the property that when they access a data item, the following accesses are likely to be to the same or nearby items.
    - This property of an algorithm is called **locality of reference**.
  - Once a memory location is accessed, the CPU cache can **prefetch** nearby memory locations. With an array, these are likely to hold nearby data items.
  - Thus, because of cache prefetching, an array can have a significant speed advantage over a Linked List, when used with an algorithm that has good locality of reference.
- ☺ With an array, iterators, pointers, and references to items can be **invalidated** by reallocation. Also, insert/remove can change the item they reference.

## More on Linked Lists Comparison With Arrays [4/4]


---

### The Moral of the Story

- Two kinds of design decisions affect the efficiency of code:
  - Choice of algorithm.
  - Choice of data structure.
- The latter often has the greater impact.

### Again:

- Use arrays when we want fast look-up/search.
- Use Linked Lists when we want fast insert & delete (by iterator).




## More on Linked Lists

### Variations — Doubly Linked List [1/3]


---

The kind of Linked List we have been discussing contains one pointer per node. Thus, it is called a **Singly Linked List**.



In a **Doubly Linked List**, each node has a data item & **two pointers**:

- A pointer to the next node.
- A pointer to the previous node.



Doubly Linked Lists often have an end-of-list pointer.

- This can be efficiently maintained, resulting in constant-time insert and remove at the end.

## More on Linked Lists

### Variations — Doubly Linked List [2/3]

---

Doubly Linked Lists have essentially all the advantages of (Singly) Linked Lists, plus some more.

- They allow efficient insert/remove at both ends of the list.
  - We can maintain an end-of-list pointer without trouble.
- They allow efficient insert-before-this-node and remove-this-node.
- They allow efficient backwards iteration.

However

- Doubly Linked Lists are a *little* slower than Singly Linked Lists.
  - Constant-time operations remain  $O(1)$ , but the constant is a little larger.
- Doubly Linked Lists still cannot do both splice and size efficiently.

The Bottom Line

- Doubly Linked Lists are generally considered to be a good basis for a **general-purpose** generic container type.
  - Singly-Linked Lists are not. Remember all those asterisks?

The C++ STL has Doubly Linked Lists: the `std::list` class template.

## More on Linked Lists

### Variations — Doubly Linked List [3/3]

	Smart Array	<b>Doubly</b> Linked List
<b>Look-up by index</b>	<b><math>O(1)</math></b>	<b><math>O(n)</math></b>
<b>Search sorted</b>	<b><math>O(\log n)</math></b>	<b><math>O(n)</math></b>
Search unsorted	$O(n)$	$O(n)$
Sort	$O(n \log n)$	$O(n \log n)$
<b>Insert @ given pos</b>	<b><math>O(n)</math></b>	<b><math>O(1)</math></b>
<b>Remove @ given pos</b>	<b><math>O(n)</math></b>	<b><math>O(1)</math></b>
<b>Splice</b>	<b><math>O(n)</math></b>	<b><math>O(1)</math></b>
<b>Insert @ beginning</b>	<b><math>O(n)</math></b>	<b><math>O(1)</math></b>
<b>Remove @ beginning</b>	<b><math>O(n)</math></b>	<b><math>O(1)</math></b>
Insert @ end	$O(1)$ or $O(n)^*$ amortized const	$O(1)$
Remove @ end	$O(1)$	$O(1)$
Traverse	$O(n)$	$O(n)$
Copy	$O(n)$	$O(n)$
Swap	$O(1)$	$O(1)$

With Doubly Linked Lists, we can get rid of most of our asterisks.

\* $O(n)$  if reallocation occurs. Otherwise,  $O(1)$ . Amortized constant time.

- Pre-allocation can help.

**Find** faster  
with an array

**Rearrange** faster  
with a Linked List

## More on Linked Lists

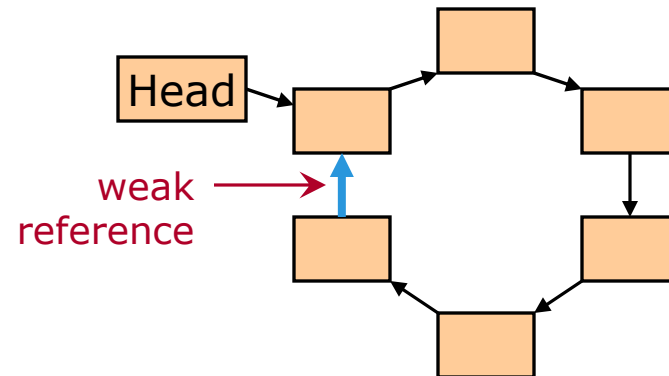
### Variations — Circular Linked List

---

Using nodes and pointers, we can arrange data structures in just about any way we want.

An interesting variation on a Linked List is a **Circular Linked List**.

- Here, we arrange our nodes in a circle.



- Ownership issues get trickier.
  - Destroy the head and ... nothing else gets destroyed?
  - One (somewhat icky) solution: use a **weak reference**.

## More on Linked Lists Implementation

---

Two approaches to implementing a Linked List:

- A Linked List package to be used by others.
- An internal-use Linked List: part of some other package, and not exposed to client code.

How would these be different?

- In particular, what classes might we define in each case?
  - In the first case, many classes: container, node, iterator.
  - In the second case, a node class, but possibly nothing else.

TO DO

- Update the internal-use Singly Linked List begun long ago, to include:
  - Exception-safety information.
  - An insert-at-beginning operation.

*Done. See `linked_list.cpp`,  
on the web page.*