

Linked lists and memory allocation



Prof. Noah Snaveley

CS1114

<http://cs1114.cs.cornell.edu>



Cornell University
Computer Science

Administrivia

- Assignment 3 has been posted
 - Due next Friday, March 6
- Prelim 1 Thursday in class
 - Review session this evening at 7:15pm, Upson 315
 - Review session tomorrow, 8:30pm?
 - Topics include: running time, sorting, selection, graphs, connected components, linked lists
 - Closed-book, closed-note

Administrivia

- Homework policy:
- You can discuss problems in general with other students
- You must write the code on your own – you may not share code

Bubble sort

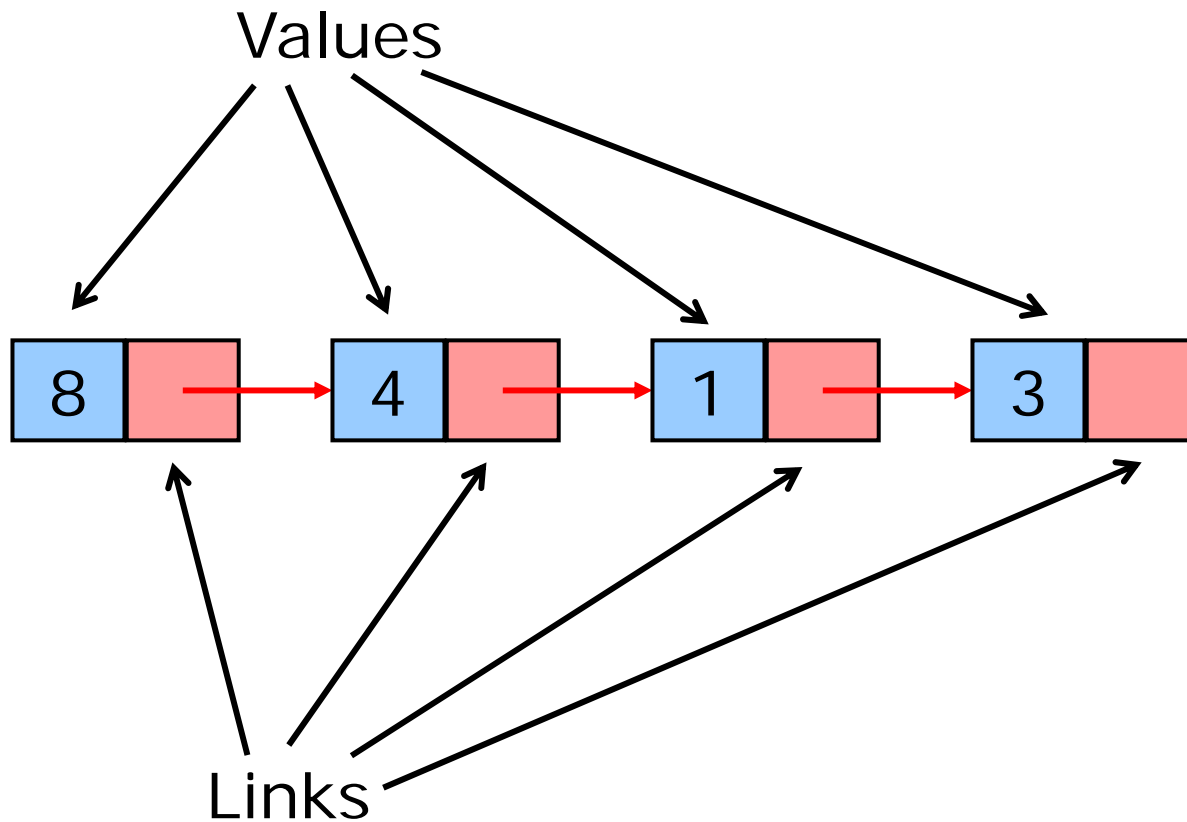
- What is the running time?
- Which is faster?
 - a) Bubble sort
 - b) Quicksort

Linked lists




- Alternative to an array
- Every element (cell) has two parts:
 1. A value (as in an array)
 2. A link (address, pointer) to the next cell

This pointer will always point to the start of a cell



Linked lists





Example






Doubly linked lists



A doubly-linked list in memory

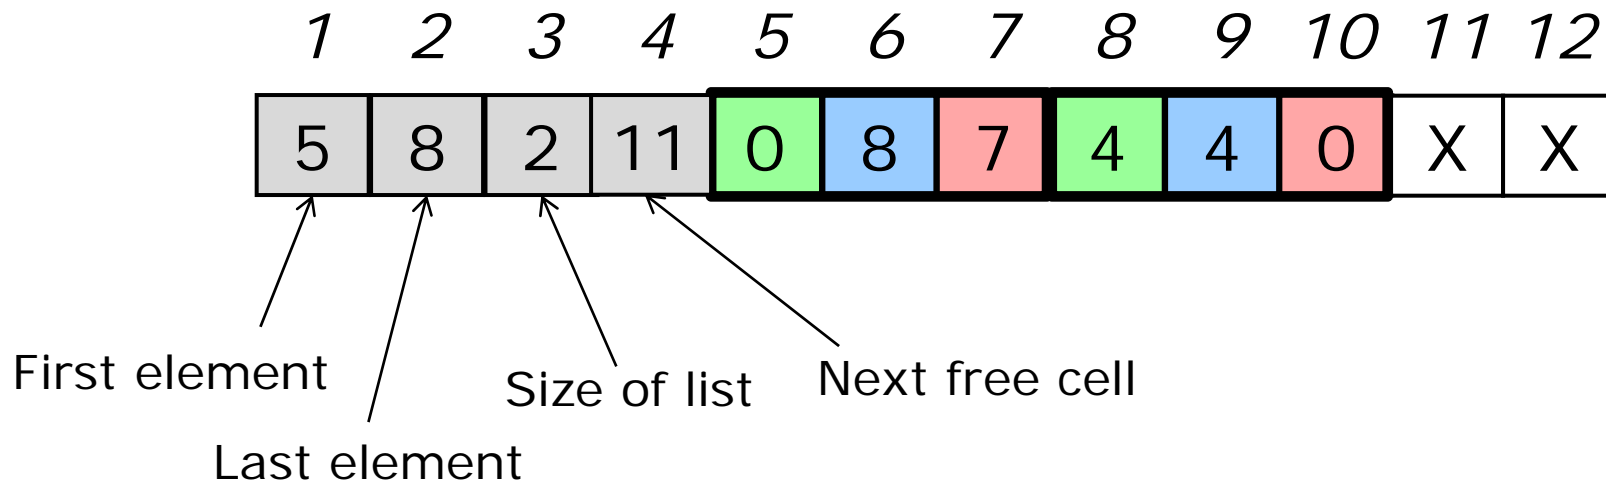


Doubly-linked list insertion




Memory allocation


- When we need a new cell, how do we know where to find it?
- We'll keep track of a "free pointer" to the next unused cell after the list




Doubly-linked list insertion



Insert a
9 at the
start



Delete
the last
element





Memory allocation

- Current strategy: when we need more storage, we just grab locations at the end
- What can go wrong?
- When we delete items from a linked list we change pointers so that the items are inaccessible
 - But they still waste space!



Memory allocation

- Strategy 1: Computer keep tracks of free space at the end
- Strategy 2: Computer keeps a linked list of free storage blocks (“freelist”)
 - For each block, stores the size and location
 - When we ask for more space, the computer finds a big enough block in the freelist
 - What if it doesn’t find one?

Maintaining a freelist



Delete
the last
element



Allocation issues

- Surprisingly important question:
 - Which block do you supply?
 - The smallest one that the users request fits into?
 - A larger one, in case the user wants to grow the array?

Memory deallocation

- How do we give the computer back a block we're finished with?
- Someone has to figure out that certain values will never be used ever ("garbage"), and should be put back on the free list
 - If this is too conservative, your program will use more and more memory ("memory leak")
 - If it's too aggressive, your program will crash ("blue screen of death")



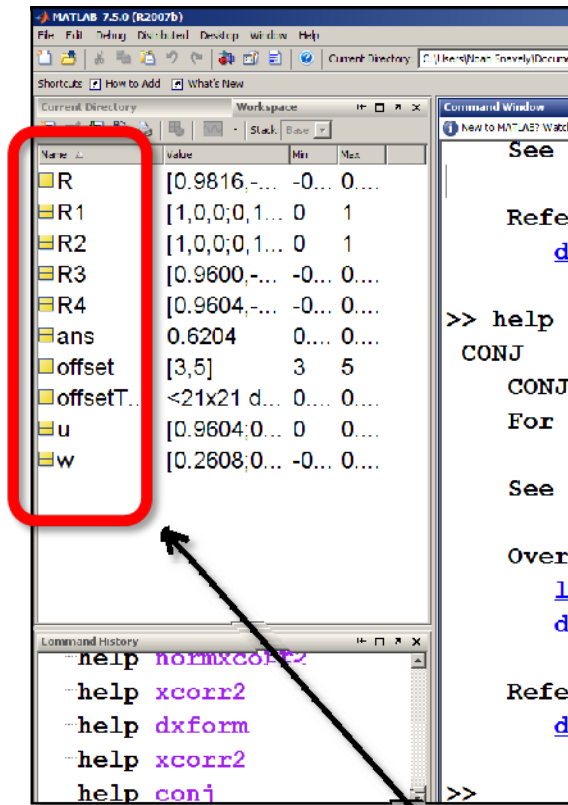
Memory deallocation

- Two basic options:
 1. Manual storage reclamation
 - Programmer has to explicitly free garbage
 - Languages: C, C++, assembler
 2. Automatic storage reclamation
 - Computer will notice that you're no longer using cells, and recycle them for you
 - Languages: Matlab, Java, C#, Scheme

Manual storage reclamation

- Programmers always ask for a block of memory of a certain size
 - In C, explicitly declare when it is free
- Desirable but complex invariants:
 1. Everything should be freed when it is no longer going to be used
 2. If we free something, we shouldn't try to use it again
 3. And, it should be freed exactly once!
 4. Minimize fragmentation

Automatic storage reclamation



- “Garbage collection”
- 1st challenge: find memory locations that are still in use by the programmer (“live”)
 1. Anything that has a name the programmer can get to (the “root set”)
 2. Anything pointed to by a live object

Root set in Matlab

Garbage collection

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
11	2	4	5	1	0	8	3	15	1	6	5	3	7	1	0
X								Y							


- Two lists, X and Y
- Which cells are live?
- Which cells are garbage?

Simple algorithm: mark-sweep

- Mark: Chase the pointers from the root set, marking everything as you go
- Sweep: Scan all of memory – everything not marked is garbage, and can go back on the free list



Mark and sweep



Root set: { X, Y }


- Mark phase
- Sweep phase

Mark and sweep

- The machine needs to be able to tell where the pointers are (we'll assume that it's up to the programmer to do that)
 - For instance, the programmer will say that the second entry in a cell is a pointer (for singly-linked list)
 - Or, for a doubly-linked list, the first and third entries in a cell are pointers

Mark and sweep

- In general, pointers may have a complex structure



How do we mark, in the general case?

When to do garbage collection?

- Option 1 (“stop the world”): Once memory is full, stop everything and run garbage collection
 - Your program will freeze while the garbage is being collected
 - Not good if you’re coding the safety monitoring system for a nuclear reactor
- Option 2 (“incremental GC”): Collect garbage in parallel with the main program
 - Needs to be careful not to step on the program

Assignment 3

- Implementing stacks and queues using linked lists
- Using DFS and BFS to find connected components
- Guiding the robot with the lightstick