Quiz 1 Review

CS50 — Fall 2008

Week of December 7, 2008

Contents

1 Helpful Study Hints

Here are a few study tips for preparing for the quiz.

- Read the Quiz 0 review section notes thoroughly! This quiz is cumulative, and you will be expected to know everything that you were supposed to know for Quiz 0! Although the quiz will prioritize material not yet tested, questions may assume an understanding of previously tested concepts. Therefore, you should review questions you missed on the first quiz to ensure those same concepts don't trip you up again!

- Practice tracing through code, and asking yourself *why* things work the way they do. If you've been doing the hacker problem sets, make sure you understand the concepts presented in the standard problem sets!

- Practice with pointers on your computer. Write a program called pointers.c and compile it for a bunch of different use cases of pointers. For example, practice passing objects by value vs. by reference into a function call, practice using * to dereference pointers, and & to get the address of objects. A good exercise would be to write a swap function using pointers.

- Look over the Makefiles that you have used and written for your early assignments. Make sure you have some idea of how they work and what they do.

- Review all the lecture notes, section notes and scribe notes. Find and make note of tricky syntax.

- The quiz is closed-book, but you're allowed to bring "one two-sided page (8.5" × 11") of notes, typed or written." Use that page to jot down details you're worried you might forget (*e.g.*, the format of a for loop in C). To be clear, it's not terribly important to have such details memorized at this point (after all, you can always look such up in the real world). But why waste time on the quiz trying to remember little things like that.

- Be familiar with the guest lectures. These are fair game, too!

- Practice writing some XHTML and PHP. Start with some simple stuff, and build up the complexity a little bit.

- Make some SQL tables and practice querying them with phpMyAdmin.

- Know all of the XHTML tags that we went over in lecture, and try to recall the bugs you frequently encountered when you were writing your own webpage in XHTML (forgetting tags, forgetting quotes, etc.).

- Review thoroughly the scribe notes, lecture videos, and code examples for JavaScript. The best way to get your head around the material is to really immerse yourself in it!

- Do all the practice questions we've written for you, both for this quiz review and the last one! Note that these section notes don't happen to include problems on JavaScript and Ajax, but they too are fair game!

2 C Topics

2.1 Functions

- Declarations: Make sure that you know how do declare and call functions. Remember that a function header always has the following format:

  ```
  <return type> <function name> (<arg1 type> <arg1 name>, <arg2 type> <arg2 name> ...)
  ```

  So, for example, the declaration of a sort function might look something like:

  ```
  bool sort(int array [], int size)
  ```

- Arguments: What happens to the arguments that you "pass" to a function? Are they copied? What can you do to ensure that an argument is not copied? We refer to this as passing by _____?

- Functional Decomposition: Functional decomposition means breaking down a program into its (logically) smallest parts. There are two main reasons that we do this:

  - Reuse of Code: Consider writing code to sort an array of integers. It would be annoying if you had to write this code every time you wanted to sort something! A much more sensible approach is to *decompose* your program by writing a sort function that takes, as an argument, the array to be sorted. In this way, any time you want to sort an array, you can just pass it to your sort function, as easy as that!

  - Readability: Functional decomposition makes code much more legible for other people reading it!

2.2 Scope and Shadowing

The **scope** of a variable is analogous to "where it can be seen" in your program. Thus, variables defined within a particular function are only visible and accessible inside of that function, after the variable declaration. Similarly, a variable defined within a for loop or a while loop will not be accessible anywhere except after where you declared the variable inside of that loop!

To generalize, each *block* of code defines a new scope, so that any variable declared in that block is accessible (after the declaration) inside of that block (including any more nested blocks within that block). You can think of blocks as the sections of code enclosed in curly braces ({ and }).

Variable **shadowing** (a term we may not have used before) means that you have two variables with the same name, declared in overlapping scopes (normally from nested blocks). We say that the inner scope declaration shadows the outer declaration because any use of that variable name in the inner scope (after the inner scope declaration) will refer only to the declaration in the inner scope. There is no way to "reach out" into the outer scope in C...so do be careful when reusing variable names.

2.3 Searching Algorithms

So far, we have studied two major searching algorithms for an array. They are *linear search* and *binary search*.

2.3.1 Linear Search

The pseudocode for linear search is as follows:

```
linear(array, element_sought):
    for each element in array
        if (element equals element_sought)
            return true
    return false
```

Linear search runs in linear time, $O(n)$. Remember that when quoting computational complexity, we always quote the worst case scenario. In the worst case of linear search, we have to compare each of the n elements of the array to what we are looking for. What is the computational complexity of this algorithm in the best case?

2.3.2 Binary Search

Binary search is more efficient, running in $O(\log n)$. However, it only works on arrays that have been sorted! A practice exercise asks you to draw up some pseudocode for the binary search algorithm.

Why does this algorithm run in $O(\log n)$? Try thinking about it as $O(\log_2 n)$ instead. Is it clearer? The algorithm eliminates half of the remaining array each pass through. And you can't do that more than $\log_2 n$ times, because $n/(2^{\log_2 n}) = 1$, and a singleton set can't be further divided. Think of this as a similar algorithm to the counting algorithm presented near the beginning of the course. Half of the remaining people in the class are counted with each pass, and the rest sit down. Thus a class containing x students will take only $\log_2 x$ passes, rounded up to the nearest integer, to count.

2.4 Asymptotic Notation

If you didn't understand what was going on in the previous section when discussing $O(n)$, then this is where you want to focus some attention! $O(n)$ is how we represent *time complexity*, also known as *computational complexity*. When we say this, what we are getting at is trying to figure out how long it will take (how much of our computational resources it will take) to solve a given problem. Some of the most common notations are as follows:

$O(1)$ - *constant time* - Given a data input of any size, the computation will always take the exact same amount of time. The following function runs in constant time, because no matter how large `array` is, the same answer is always returned:

```
int func(int array[]) {
    return 2;
}
```

$O(n)$ - *linear time* - Given a data input of any size, the computation will run in an amount of time proportional to the size of the list, linearly. This means that given a data input of size k, adding one element to the data input makes the computation take one unit of time longer to run. An example of this is the linear search algorithm, outlined above.

$O(\log n)$ - *logarithmic time* - Given a data input of size n, the computation runs in $\log_2 n$. Examples of this include the binary search algorithm, outlined above, as well as the counting algorithm discussed in Week 4.

$O(n^2)$ - *quadratic time* - Given a data input of any size, the computation will run in an amount of time proportional to the size of the list, squared. This means that given a data input of size k, adding one element

to the data input makes the computation take $(k+1)^2 - k^2 = (2k+1)$ units of time longer to run. Examples of this include bubble sort and selection sort.

Some of the other useful time complexities to know are *supralinear time*, $O(n \log n)$ [an example of this is merge sort]; *polynomial time*, $O(n^c)$; *exponential time*, $O(c^n)$; and *factorial time*, $O(n!)$. All of these are with respect to a data input of size n.

2.5   Recursion

A very important topic to remember for this quiz is the idea of **recursion**. Remember that an algorithm is recursive if it (1) breaks a problem down into smaller versions of itself, and (2) has one or more base cases to which we know the solution. The practical result of this is a function that calls itself.

A good example to demonstrate recursion is any algorithm that computes a summation or a factorial. Look at the following code that gets the sum of all the integers from 1 to `max` using a recursive approach:

```
int sum_recursive(int max) {
    if(max < 1) {
        printf("Illegal argument. Sum cannot be computed.\n");
        return 0;
    }

    if(max == 1)
        return 1;
    else
        return n + sum_recursive(max - 1);
}
```

A practice example exercise asks you to convert this back to an iterative algorithm. Remember that the two components of a recursive algorithm are (1) the recursive process and (2) the base case. We can see from the algorithm the base case: it is when `max` is equal to 1; that is when the algorithm (the for loop) stops doing work and moves on. Intuitively, we know that this equates to the sum of all the numbers from 1 to 1, which we know to be 1. This is the simplest case of the algorithm. The recursive process comes from noting that, for example:

```
sum_iterative(4) = 4 + 3 + 2 + 1
sum_iterative(3) =     3 + 2 + 1
sum_iterative(2) =         2 + 1

or...

sum_iterative(2) =         2 + 1
sum_iterative(3) =     3 + sum_iterative(2)
sum_iterative(4) = 4 + sum_iterative(3)
```

2.6   The Stack

Another important thing to consider when discussing recursion is **the stack**. Think of the stack as a pile of function "frames". The last function you called has a frame on top of the stack, while the first function you called has a frame on the bottom of the stack. Applying this to the `sum_recursive()` function, above: If we call `sum_recursive(4)` from `main()`, what does the stack look like just before the base case returns its value?

```
sum_recursive(1)    --  called by sum_recursive(2)

sum_recursive(2)    --  called by sum_recursive(3)

sum_recursive(3)    --  called by sum_recursive(4)

sum_recursive(4)    --  called by main()

    main()
```

Why is this the case?

2.7   Sorting

Again, we have focused extensively on some sorting algorithms thus far. Three in particular are going to be important for you to know: *bubble sort*, *selection sort*, and *merge sort*. Take a look at the section notes that describe these algorithms in detail.

Some of the take home points are these:

- Bubble sort and selection sort run, in the worst case, in $O(n^2)$. Read through the algorithms that these sorts use, and make sure you understand why this is the time complexity for each.

- Merge sort runs in $O(n \log n)$. Note that merge sort is a naturally recursive algorithm. Make sure you understand why we describe it as a recursive algorithm, and why the time complexity is as it is.

There are other sorts that you may want to know or use on the Quiz. These include insertion sort, quicksort, cocktail sort (aka bidirectional bubble sort), etc. Knowing the time complexity and general algorithms behind these will be sufficient preparation for you.

2.8   Pointers

If you remember nothing else about pointers, remember this: **Pointers are Addresses!**
Pointers are variables, just like any other variables, but instead of holding (for example) a character, they hold an address.

Syntax   Remember that you declare pointers by inserting an asterisk (∗) in front of the variable name.

```
1       int *iptr;
2       int* iptr1, var2;
```

Note the in the previous example, the first statement declares one pointer, and the second statement declares one pointer and a normal variable. Just because I put the asterisk near int does NOT mean that all variables defined on that line are integer pointers.

Remember that pointers store *addresses* of variables rather than the data directly. If you have an address, you can use the ∗ operator to get to the value. To find the address of a variable, use the & operator. See below for an example.

```
1       int var;
2       int *iptr;
```

```
3          iptr = &var;
4          *iptr = 4; /* this will set the var = 4 */
5          iptr = 4; /* invalid, will compile, but definitely not what you want */
```

Remember that a pointer can only hold the address of a certain type of variable, specified when you declared it. In the above example, iptr can only hold the address of an integer, and not any other type.

Passing by Reference Before pointers, we could only pass arguments to function *by value*. This meant that when we called a function foo with an argument x, the value of x would be copied to a new local variable in foo. No matter what happened to this local variable, it did not change x in the calling function. In other words, we passed only the *value* of x to foo, but not x itself.

With pointers, we can now give functions access to the *original* variable. We do this by passing *by reference*. We pass the address of the original x (&x) to foo, and foo stores it in a pointer. Now, when foo dereferences that pointer, it grabs a hold of x.
Why would we ever want to do this? When x is very large, like a huge array or a big structure, it takes up a lot of memory to keep making new copies of it in each new function. It is much more efficient to have one copy of it, and just tell all the functions where to find it!

Arrays and Pointers Arrays and pointers are very interrelated. In memory, an array is a contiguous block of memory. An array of 10 integers is, in memory, a block of $(10 * \texttt{sizeof(int)})$ bytes. As the array name is implicitly converted into a pointer to the first element of the array, the array name can be thought of as a pointer to the beginning of this block.
This means we can treat the array name as the address of the array, and assign it to a pointer:

```
1          int arr[10] = {0,1,2,3,4,5,6,7,8,9};
2          int *ptr;
3          ptr = arr;  /* ptr holds the address of the beginning of the array */
```

The indexing operators ('[' and ']') are interrelated with pointer arithmetic. The following boolean expressions all evaluate to TRUE:

```
(     ptr[3]  ==  arr[3]    )
( *(ptr + 3)  ==  arr[3]    )
(    arr + 5  ==  &(ptr[5]) )
(   *(++ptr)  ==  arr[1]    )
```

2.9   Dynamic Memory

malloc malloc is a function, declared in stdlib.h, that returns memory allocated off the heap. When we declare a pointer, it is *uninitialized* - it does not point to any memory that our program owns. If we already have the memory that we want our pointer to point to, then we can use the & operator to get the address. However, if we don't have that memory yet, we use malloc to get it dynamically. If there was an error in allocating the memory, malloc returns NULL, so we always must check for this return value before proceeding.
This is useful if, for example, we want an array, but we don't know how big it should be until the program is already running.
    malloc allocates memory from the *heap*. The heap is the other main section of memory for programs, along with the stack. The heap is a large pool of memory that we may dip into. While the stack is a tightly policed city park, where local variables and function stack frames are created and wiped out according to strict rules, the heap is a country pasture that is unsupervised. Once you grab some memory from the heap, it sticks around until you release it - you can keep it forever, as long as your program is running. This is great, because as long as you keep the address of your memory in the heap, you can keep using it no matter

7

what function you are in. The downside is you *must* keep track of its address. When you are done with memory on the heap, you must free it with the function free. It is a serious error to not free all the memory you allocate.
All memory allocated with malloc must be freed with free.

2.10   User-defined Types

2.10.1   typedef

The keyword typedef allows the naming or renaming of types in C.
Its syntax is:

```
typedef <old type> <new name>;
```

2.10.2   enum

Enums are types of variables that can hold one of an enumerated list of possible values. When you enumerate the possible values, they are assigned integer equivalents, starting with 0 for the first one, 1 for the second one, and so on.

```
1          enum quad_house_t {CURRIER, PFOHO, CABOT);
2
3          enum quad_house_t johns_house = CABOT;
```

2.10.3   struct

Structs are variable types that combine more than one variable.

```
1          struct student_t {
2                  char *name;
3                  int year;
4                  int age;
5          };
```

    If you have a pointer to a struct, then you may use the -> operator to access member variables. You may also dereference the pointer and then use the . (dot) operator. For instance, the following piece of code accesses members of the struct in three different ways.

```
1          struct student_t student1;
2
3          int main(int argv, char * argv[]) {
4              struct student_t *studentptr = &student1;
5              student1.name = "Bobby";
6              studentptr->year = 1;
7              (*studentptr).age = 19;
8              return 0;
9          }
```

2.11   File I/O

C programs can interact with other files using the file input/output functions defined in stdio.h. An excellent resource for information about all the functions in stdio.h is online, at http://www.cplusplus.com/reference/clibrary/cstdio/.

8

Within our programs, we interact with files using a special variable type, a `FILE*` (file pointer). We get one of these by calling `fopen`, which returns a `FILE*`. Two key file I/O functions are `fread` and `fwrite`, which behave a bit differently than functions we are used to. The return type of `fread` is `size_t` so where does it put the data that it read from the file? It puts it in a location that we specify when calling the function, by passing it a pointer. Take a look at `fread`'s function header:

```
size_t fread ( void * ptr, size_t size, size_t count, FILE * stream );
```

We have to pass in the address of the location for `fread`. We pass this in as the first argument, called `ptr`. The return value is the number of bytes read.

2.12   Digital Forensics

In Problem Set 5, we learned about what happens to files when they are "deleted" from a hard drive. Remember, the reference to them in the drive's file system is erased, but the data is not. This allows someone to go back to that drive, and read all the data, as we did in the problem set.

How did we recognize different files, when there was no directory structure? We looked at the file *headers*, which contain metadata about the file, rather than the actual file data itself.

Remember the different ways to truly erase a hard drive. We could physically destroy it, or degauss it, which are the only totally secure ways of data deletion. There are methods that don't render the hard drive useless, but none of these are 100% safe. Make sure you understand why just overwriting data doesn't prevent recovery of the old data.

2.13   Algorithms of Composition

A common problem in computer science is that of finding the best way to make something out of the smallest number of parts possible. A good example of this is the "making change" problem: Given that I have unlimited supplies of coins in the following denominations: 2c, 5c, 10c, and 25c, what is the smallest number of coins needed for me to make up x cents?

The Greedy Algorithm   A greedy algorithm for making x employs the following rule:

Pick the largest denomination possible (i.e. still smaller than x). Subtract this value from x; let this value be the new x. Repeat until x = 0.

For example, say that x = 70c. If I were using a greedy algorithm, I would pick the following coins in this order:
- 25c (now we need 45c)
- 25c (now we need 20c)
- 10c (now we need 10c)
- 10c(now we're done!)

The Exhaustive Search   A greedy algorithm doesn't work in all cases. In fact, consider what would happen in the case above if we were trying to make 28c instead of 70? We would get to the point where we had 1c left - and no way to make up that 1c! It is also important to remember that even if a greedy algorithm *does* yield a solution, it might not be the optimal solution. To find the optimal solution, we must employ an exhaustive search.
    The idea behind an exhaustive search is that we find all possible combinations of denominations that make up our total goal. We then compare these, and use the one with the fewest units as our optimal solution.

For example, with denominations 2c, 5c, 10c and 25c, I can make up 9c as follows
1: {5c, 2c, 2c}
2: {5c, 1c, 1c, 2c}
3: {5c, 1c, 1c, 1c, 1c}
Clearly the optimal solution is solution 1. In this case, this is also the solution that a greedy algorithm would have returned.

2.14   Linked Lists

Be sure to consider answers to all of the following questions when studying.

- What is our motivation for the use of linked lists? How do linked lists compare to arrays? How do they differ? What is one advantage provided by linked lists that is not provided by arrays? In terms of memory, how do arrays and linked lists differ?

- How do we represent singly linked lists in C? What are the general algorithms for inserting, deleting, and finding elements in a singly linked-list?

- What is the time required to do an insertion or deletion? What is the time required to perform a lookup? How do these differ from arrays?

- How is the struct definition for a doubly-linked list different from that of a singly-linked list? When it is it advantageous to use a doubly-linked list instead of a singly-linked list?

2.15   Binary Trees

Past section notes will prove a valuable resource for you when studying this topic.

- How do we go about implementing a binary tree in C? What is its structure definition? How do we declare a new tree?

- What distinguishes a binary search tree from ordinary binary trees? How do we insert elements into a binary search tree? How do we search the binary search tree? What is the runtime of this search algorithm?

- What are the two distinguishing properties of heaps? How do you heapify a non-heap? What is the runtime of heapification and why? How do we use heaps to implement heap sort? What is the runtime of heap sort and why?

2.16   Hash Tables

- Although hash tables, when implemented with separate chaining, can have asymptotically the same lookup time as linked lists searched linearly, in practice they are usually much more efficient. Why might that be?

- How do we typically represent a hash table in C? How would you declare an instance of a hash table?

- Why are collisions a problem with hash tables but not the other data structures we have studied? What are the two ways to deal with collisions in a hash table?

- What is the complexity of the lookup operation? What is the complexity of the insert and delete operations?

2.17 Miscellaneous C topics

If you feel confident with all of the above topics, and the big topics in the web programming section, then feel free to dive into these:

- Huffman Coding
- Bitwise Operators
- The C Preprocessor
- Compilers

3 Web Programming

We **strongly** encourage you to review the Section Notes for weeks nine and ten, as they contain comprehensive detail into JavaScript and PHP/SQL. Also, review Keito's supersection and walkthroughs, and the scribe notes, as they will be very useful resources when you are brushing up on this material.

3.1 TCP/IP

TCP/IP stands for Transmission Control Protocol / Internet Protocol. It is a set of rules that define how most electronic devices, namely computers, should be connected to the Internet, and how data should be transferred between electronic devices. TCP and IP refer to different protocols which handle different aspects of this communication.

Take home points:

- Every computer has an IP address that uniquely identifies it on the Internet
- This unique address lets you send data (in the form of IP packets) across the internet
- Routers perform the physical task of figuring out where to send your data to get it to the correct IP address
- IP is connection-less in that it doesn't specify the route that data will take
- TCP handles packets once they arrive at a computer, determining what application they go to.

3.2 HTML and XHTML

Here are some of the important things to remember (Section Notes 9):

- Basic components of page structure: `<html>`, `<head>`, `<body>` tags
- Basic tags: Links, images, headings
- Tables
- Forms
- GET vs. POST
- Differences between XHTML and HTML

3.3 CSS

Cascading Style Sheets provide a way to override basic XHTML tags.

- Basic syntax
  - Selectors
  - Declarations
  - Arguments to
- CSS Classes
- Embedding CSS in XHTML

3.4 PHP

- What advantages does PHP have over plain XHTML?
- `$_GET`
- `$_POST`
- `$_SESSION`
- Connecting to SQL databases

3.5 SQL

SQL (Structured Query Language) provides an interface to server-side databases which you can create.

- SELECT
  - FROM
  - WHERE
  - ORDER BY
- UPDATE
- INSERT
- DELETE

3.6 JavaScript

From here on out, reference Section Notes 10.

- Variables and lack of type
- Arrays
  - Differences between C arrays and JavaScript arrays
- `+, ==, ===`
- Determining the type of a variable
- `switch`

- `for...in`, and syntactic differences between that and a `for` loop in C

- Functions
  - What does it mean to be a first-class object?
  - Event handlers

- Objects
  - Properties
  - Methods
  - Constructors

- The Document Object Model [DOM]
  - Properties of DOM
  - Methods of DOM

3.7 AJAX

- What is AJAX used for?

- XMLHttpRequest
  - Attributes of the class
  - Methods of the class

- Integrating AJAX, XHTML, and PHP

- Google Maps and the Google Maps functions

4 Problems

1. Declare a function that takes an integer as an input and returns true if the number is even and false if the number is odd.

2. Write the function you declared in Question 1.

3. Modify the function you declared and wrote in Questions 1 and 2 to instead accept a pointer to an integer instead of the integer itself.

4. Consider the below function, which deals with variable scope:

```
1    #include <stdio.h>
2
3    int x = 10;
4
5    int main() {
6        printf("The result is %d and x is %d here\n", f(4), x);
7    }
8
9    int f(int x){
10       return x * x;
11   }
```

What is printed at the end of this function?

5. Write a pseudocode algorithm for binary search.

6. How many steps will it take for a binary search algorithm to look through the array
4 6 9 11 14 18 21 25 29 32 37 to find the element 32?

7. What does the following code snippet do:

```
1    void mystery(char *foo) {
2        if(*foo == '\0')
3            printf("\n");
4        else {
5            printf("%c", *foo);
6            mystery(foo+1);
7        }
8    }
```

8. State the time complexity of the algorithm given in Question 7 using asymptotic notation.

9. Compare the function in Question 7 to the following one:

```
1    void *new_mystery(char *foo) {
2        if(*foo == '\0') {
3            printf("\n");
4            return;
5        }
6        else {
7            new_mystery (foo + 1);
8            printf ("%c", *foo);
9            return;
10       }
11   }
```

Do they have the same effect?

10. Bubble sort and bi-directional bubble sort both run in $O(n^2)$ time, even though when run side by side, the latter appears to move quicker. Explain why both run in the same time complexity.

11. Write an algorithm that prompts a user to repeatedly enter a series of `doubles` until they type in some sentinel value, and return to them the average of these values. Describe the efficiency of your algorithm using big O notation.

12. Let us consider the following array of integers:
5 8 2 1 6 7 3 4
and the uni-directional bubble sort algorithm, where we bubble the largest values to the right on each pass through, and then return to the front of the array to do the same, until the array is sorted. If we use this algorithm, what does the array look like after (a) one pass? (b) two passes? (c) three passes?

13. Write a recursive function that computes the n^{th} Fibonacci number. (The sequence begins 0, 1, 1, 2, 3, 5, 8...)

14. Imagine we have the following functions declared in a C file:

```
1    void another_level() {
2        another_level();
3    }
4
5    void main() {
6        another_level();
7    }
```

If we run the program, will it ever terminate? If so, why? If not, why not?

15. Say my program looks like this:

```
1        ...
2        int i;
3        int *ptr;
4        i = 5;
5        /* ?? */
```

What line should replace the question marks, so that `ptr` holds the address of `i`?

16. If we have this line in our program:

```
1        int array[] = {1,2,3,4};
```

What is the value of the following expression?

```
        *(&(*(array + 1)))
```

17. How do we declare, dynamically, an array of 40 `doubles`?

18. Point out at least four things wrong with the program that appears below:

```
1   #include <stdio.h>
2
3   #define NUM_INTS 22;
4
5   int main() {
6       int i;
7       int *myInts = malloc(NUM_INTS * sizeof(int))
8
9       for(i = 0; i <= NUM_INTS; ++i)
10          myInts[i] = (i*i);
11
12      for(j = NUM_INTS; j > 0; j--)
13          printf("%d\n", myInts[i-1]);
14
15      return 0;
16  }
```

19. Define a structure for a type that can hold the number of students enrolled, the name of the instructor, and the CUE rating (just the overall score).
    Then, demonstrate how to declare an instance of this structure.

20. If we had an instance of the structure declared in Question 19 , how would we access the "instructor" field?

21. If we had a pointer to an instance of the structure declared in Question 19, how would we access the "number of students" field?

22. If you are selling your hard drive on eBay, why might it be a good idea to do a complete format, rather than a "quick format"? (Even if you don't know the precise difference between the two, try to reason based on your knowledge of what happened with David's Flash card!)

15

23. Demonstrate how to open an ASCII file named "input.txt" for reading.

24. Demonstrate how to open an ASCII file named "output.txt" for writing.

25. In `whodunit.c`, from Problem Set 5, what was the underlying representation of our `BYTE` type? That is, you can declare the BYTE type as:
    `typedef _____ BYTE;`

26. How would you define a new enumerated type called `months` that, understandably enough, enumerates all of the months of the year.

27. What will the following code segment return?

```
1        #include <stdio.h>
2        int x = 10;
3        int y = 20;
4
5        int main() {
6            return f(x + 4, y);
7        }
8
9        int f(int x, int z){
10           y +=x;
11           x +=y;
12           return x + z;
13       }
```

Questions 28-29. Take a look at the following problem and answer the questions below:

How do I dispense $1.43, using as few coins as possible, if I have the following coin denominations: 1, 5, 10, 20, 25, 100, and 200 cents

28. If I use the greedy algorithm?

29. If I use an exhaustive search algorithm?

Questions 30-33. Take a look at the following code and answer the questions below that relate to file I/O:

```
1   #include <stdio.h>
2
3   int main() {
4     FILE *sourceFile;
5     FILE *destinationFile;
6     char buffer [1000];
7     int n;
8
9     sourceFile = fopen("file.c", "r");
10    destinationFile = fopen("file2.c", "w");
11
12    if(sourceFile==NULL) {
13      printf("Error: can't access file.c.\n");
14      return 1;
15    }
16    else if(destinationFile==NULL) {
17      printf("Error: can't create file for writing.\n");
18      return 1;
19    }
```

16

```
20      else {
21        n = fread(buffer, 1, 1000, sourceFile); /* grab all the text */
22        fwrite(buffer, 1, n, destinationFile);  /* put it in file2.c */
23        fclose(sourceFile);
24        fclose(destinationFile);
25
26        destinationFile = fopen("file2.c", "r");
27        n = fread(buffer, 1, 1000, destinationFile); /* read file2.c   */
28        printf("%s", buffer);                        /* display it all */
29
30        fclose(destinationFile);
31        return 0;
32      }
33    }
```

30. On lines 9 and 10, what do the "`r`" and "`w`", respectively, do?

31. On lines 9 and 10, what do the calls to `fopen(...)` return?

32. Explain the flow of information in lines 21 and 22.

33. In line 26, note that `destinationFile` is being opened again. Why do this instead of just keeping it open in the first place?

34. Write a function `delete_list()` that takes a linked list, deallocates all of its memory, and sets the head to NULL.

35. Write a function `reverse_list()` to reverse a singly-linked list.

36. The following is an excerpt from the function `insert()`, which attempts to insert a node called newptr (already instantiated and initialized) into a linked list pointed to by first. This linked list is composed of nodes that point to student structures. But while debugging, we find that this code frequently hangs at the command line after the first few insertions. Why? We can assume there are no duplicates in the list.

```
1         // check for empty list
2         if (first == NULL)
3             first = newptr;
4
5         // else check if student belongs at list's head
6         else if (newptr->student->id < first->student->id)
7         {
8             newptr->next = first;
9             first = newptr;
10        }
11
12        // else try to insert student in middle or tail
13        else
14        {
15            node * predptr = first;
16            while (TRUE)
17            {
18
19                // check for insertion at tail
20                if (predptr->next == NULL)
21                {
22                    predptr->next = newptr;
```

```
23                    break;
24                }
25
26                // check for insertion in middle
27                else if (predptr->next->student->id > newptr->student->id)
28                {
29                    predptr->next = newptr;
30                    newptr->next = predptr->next;
31                    break;
32                }
33
34                // update pointer
35                predptr = predptr->next;
36            }
37    }
```

37. The following function is intended to locate the node containing an integer specified by the user within a linked list pointed to by the global pointer variable first. This code wont compile. Why?

```
1  void
2  find()
3  {
4      // prompt user for number
5      printf("Number to find: ");
6      int n = GetInt();
7
8      // get list's first node
9      node * ptr = first;
10
11     // try to find number
12     while (ptr != NULL)
13     {
14         if (ptr.n == n)
15         {
16             printf("\nFound %d\n", n);
17             sleep(1);
18             break;
19         }
20         ptr = ptr.next;
21     }
22 }
```

38. Implement a function `void delete(node *n)` to remove a node from a doubly-linked list. You can assume pointers to the head and tail of the list are global. Remember to account for the cases in which the node you are deleting is the first or last node!

39. Given any binary tree, write a function to count the number of nodes in the tree.

40. Given a binary search tree of integers, the following function is supposed to print out the numbers in sorted, increasing order. But when we run it, the numbers do not appear to be in order at all! What is wrong with it?

```
1  void printTree(treenode * node) {
2      if (node == NULL) return;
```

Figure 1: Figure for Problem 41

Figure 2: Figure for Problem 42

```
3
4          printTree(node->left);
5          printTree(node->right);
6          printf("%d ", node->data);
7    }
```

41. Draw the steps the heapification algorithm would follow in order to heapify the non-heap in Figure 1.

42. One of the advantages of a heap is that, due to its completeness and heap-order properties, we can represent it as an array. For instance, the heap in Figure 2 would be represented by the array 60 35 41 15 22:

    As you can see, the children of `A[i]` are `A[2i+1]` and `A[2i+2]`. Write a function `printAsArray` to take a heap stored with pointers and print what the corresponding array would look like.

43. Assuming a completely random hash function which can take on m different values, what is the probability of a collision when we enter $n < m$ elements in the table?

44. Write a function to count the number of strings stored in a hash table. Assume the hash table is an array of size `LENGTH` and that it contains pointers to linked lists. The pointers are of type `node*`.

45. In the worst case, separate chaining lookup can take n steps. Give an example of when this would happen.

46. Think about the BIG BOARD from Problem Set 6. Many students who implemented hash tables on Problem Set 6 got different runtimes from one another. Why?

47. We can reduce the likelihood of collisions by increasing the size of the hash table. What can be a drawback of making the hash table so large that the likelihood of collisions is almost zero?

19

48. Consider the code below, which prompts a user for strings and enters them in a hash table. Draw a diagram of the hash table that results after a user inputs the following strings in this order: lemon, clementine, peach, Cherry, apple, kiwi, banana.

```
1    typedef struct _list
2    {
3          char *word;
4          struct _list *next;
5    } list;
6
7    int main() {
8          list * hashtable[10];
9          char * word = (char *)malloc(MAX_LENGTH * sizeof(char));
10
11         for (int i = 0; i < 10; i++)
12               hashtable[i] = NULL;
13
14         int total;
15         int i;
16
17         while(TRUE)
18         {
19               printf("Please enter a string");
20               word = GetString();
21               if (strcmp(word, "") == 0)
22                     break;
23
24               hashvalue = word[0] % 10;
25               list * newptr;
26
27               if((newptr = (list *)malloc(sizeof(list))) == NULL)
28                     return 1;
29
30               newptr->next = hashtable[hashvalue];
31
32               hashtable[hashvalue] = newptr;
33               newptr->word = word;
34         }
35         return 0;
36   }
```

49. The following code is a security nightmare. Identify the problems and rewrite the code to have the same functionality, but with greater security.

```
1    #include <stdio.h>
2    int main(int argc, char **argv) {
3          char buf[256], buf2[256];
4          printf("Please enter a string\n");
5          scanf("%s", buf);
6          strcpy(buf2, argv[1]);
7          printf("The first command line argument you supplied was \n");
8          printf(buf2);
9          printf(" the string you entered was ");
10         printf(buf);
11         return 0;
12   }
```

20

username	password	fullname
cs50stud	h4ck3r	CS50 Stud
malan	djmftl!	David Malan

Table 1: Table for Problem 52

50. David has put the following HTML code in his webpage:

```
<IMG SRC=blueflax.jpg ALT="Blue Fax (Linumlewisii)" WIDTH=300 HEIGHT=175 ALIGN=left>
```

Rewrite this as valid XHTML.

51. David is debating where to place the following line of code in his new valid XHTML web page:

```
<link href="css/styles.css" rel="stylesheet" type="text/css" />
```

Where should he put it?

(a) Within `<body>` tags
(b) Within `<a>` tags
(c) Within `<head>` tags
(d) Within `<rel>` tags

52. Table 1 is in your SQL database: Write the SQL command to do the following:

(a) Grab malan's password
(b) Change malan's password to "n3rd"

53. Say we want to recreate the same table (same exact rows and columns) using HTML. Complete the following code:

```
<? $result = mysql_query("SELECT * FROM users"); ?>
<tr><td>username</td><td>password</td><td>fullname</td></tr>
<? while ($row = mysql_fetch_array($result)) { ?>
<tr>
<! TODO! Replace this code! -->
</tr>
<? } ?>
```

54. What does `chmod` do (feel free to consult the man page)? When we run `ls -l` on a directory, we may see a line such as the following:

```
-rwx------ 1 username student 1000 2008-12-04 20:03 buy.php
```

That means that buy.php is readable, writable, and executable by you and no one else. If (and dont actually do this) we wanted to modify `buy.php` so that it was

(a) Readable and writable, but not executable, by you.
(b) Readable and executable, but not writable, by your group.
(c) Executable, but neither readable nor writable, by the world.

What octal numbers would replace the question marks in `chmod ???  buy.php`? If we ran `ls -l` again afterwards, what would replace the question marks on the next line?

```
-????????? 1 username student 1000 2008-12-04 20:03 buy.php
```

55. TCP/IP is a "language" that is spoken by computers as a means of transferring data over a network connection. The data though, is not transferred raw. What is the term for these "blocks" of data that are transferred?

Questions 56-57. Consider the following XHTML snippet:

```
1  <form action="register.php" method="post">
2  <table border="0" style="text-align: left;">
3  <tr>
4  <td>Username:</td><td><input name="username" type="text" /></td>
5  </tr>
6  <tr>
7  <td>Password:</td><td><input name="password" type="text" /></td>
8  </tr>
9  </table>
10 <input type="submit" value="Register!" />
11 </form>
```

56. If we wanted to store the username in some variable `$username`, what line of PHP code would we write in `register.php` to accomplish this?

57. The Password field of this form has an inherent weakness. How would you improve that line of code to make the entered password a little more secure?