Question 1. (20 points) A bit of C++ coding. As we know, a web page is formed from a mix of tags like <title> or </p> that control formatting and layout, and plain text content that is displayed as-is. One of your colleagues is working on a program to discover how frequently different tags are used and has asked for your help in writing a key function to do the counting. The input to this function is a C++ list<string> that contains the individual words and tags from a set of web pages. The function is required to allocate a new map<string,int> on the heap and store in it a collection of (string,int) pairs, where each string is a tag and the corresponding int is how often it appears in the input list.   The function should return a pointer to the newly-allocated map.

Example: Suppose the input list contains the following strings:

```
"<title>" "fun" "stuff" "</title>" "<p>" "this" "is"
"interesting" "</p>" "<p>" "and" "so" "is" "<this" "</p>"
```

Then the output map produced by the function for this input list should contain the following (string, count) pairs in any order; the ordering is not specified:

```
  ("<title>", 1), ("</title>", 1), ("<p>", 2), ("</p>", 2)
```

You may assume that all strings in the input list have at least one character, and that strings do not have extra leading or trailing whitespace that needs to be deleted.

Note: A tag is a string that begins with "<" and ends with ">". All other strings should be ignored by the function.

Note: The input words might be from a well-formed, properly nested web document, or they might be complete nonsense. Your function should not attempt to analyze the input beyond determining for each input string whether or not it is a tag and, if it is, count it.

Write your answer on the next page. You can remove this page if you wish.

Question 1. (cont.) Write your implementation below. The function heading is given for you, and you should assume that all necessary headers have already been #included. Hint: you probably won't need nearly all of this space.

```
// Return a new heap-allocated map that stores (string,int)
// pairs, where the strings are html tags (e.g., "<thing>")
// and each int is the number of times that string occurs
// in the input list of words. The list of words may contain
// other strings that are not tags and are not counted.
map<string, int> * tagmap(const list<string> & words) {

  map<string, int> * freq = new map<string, int>();

  for (auto w: words) {

    if (w[0] == '<' && w[w.length()-1] == '>') {

      (*freq)[w] += 1;

    }

  }

  return freq;

}
```

Notes: There are many possible ways to implement this function and correct solutions received full credit. `w.size()` can be used instead of `w.length()` – they do the same thing. It is possible to use `auto` for the type of the map variable. We included the explicit type in the solution above to provide a few additional details.

Question 2. (12 points) Smart pointers. The new summer intern, A. Hacker, has decided that C++ smart pointers are the solution to all memory management problems and has been going through the company's code repository replacing all pointers (things with type `T*`) with C++ shared pointers (with type `shared_ptr<T>`).

One of the packages that Hacker has changed is one that uses a double-linked list. The original code had nodes that were defined like this:

```
class IntNode {
  int val;          // payload
  IntNode *next;    // next/prev pointers
  IntNode *prev;
};
```

Hacker has replaced that definition with this one:

```
class IntNode {
  int val;                   // payload
  shared_ptr<IntNode> next;  // next/prev ptrs
  shared_ptr<IntNode> prev;
};
```

(a) (6 points) One of the other programmers tells Hacker that this can't work. Even though it uses shared pointers everywhere, code that uses the new nodes will still have memory leaks. Is the other programmer right? How is this possible? (Be brief)

Yes. If this node structure is used, every double-linked list with more than one node will contain cycles with nodes whose reference counts will never decrease to 0, and will never be automatically deleted by the `shared_ptrs`.

(b) (6 points) After listening to the objections, Hacker proposes replacing all of the `shared_ptr<IntNode>` declarations with `weak_ptr<IntNode>`. With this change, will memory management work properly with no leaks or other bugs? Be brief.

No, that won't work. `weak_ptrs` do not have any interaction with the automatic deletion done by `shared_ptrs` or `unique_ptrs`. These nodes will behave just like nodes with regular pointers and will not be reclaimed unless they are deleted explicitly.

Question 3. (16 points) C++ classes. Consider the following program, which compiles and links successfully. (What happens after that is something we'll get to later. ☺)

```cpp
#include <iostream>
using namespace std;

class Base {
public:
  Base() {
    cout << "Base constructor" << endl;
    ia_ = new int[5];
  }
  virtual ~Base() {
    cout << "Base destructor" << endl;
    delete[] ia_;
  }
  Base& operator=(const Base& rhs) {
    if (this != &rhs) {
      delete[] ia_;
      ia_ = rhs.ia_;
      cout << "Base assignment" <<  endl;
    }
    return *this;
  }
private:
  int *ia_;
};

class Derive : public Base {
public:
  Derive() {
    ja_ = new int[5];
    cout << "Derive constructor" << endl;
  }
  virtual ~Derive() {
    cout << "Derive destructor" << endl;
    delete[] ja_;
  }
private:
  int *ja_;
};

int main() {
  Base b1;
  Derive d1;
  b1 = d1;
  return 0;
}
```

(Question continued on next page – you may remove this page if you wish.)

Question 3. (cont.) (a) (8 points) What does this program print when it is executed?

(Reminder/hint: when an object of a derived class is constructed, the base class constructor for that object executes before the derived class constructor. When the object is deleted, the destructors run in the reverse order – derived class destructor first.)

Base constructor
Base constructor
Derive constructor
Base assignment
Derive destructor
Base destructor
Base destructor

(b) (8 points) Unfortunately, after the program finishes printing the output you described in your answer to part (a), it crashes and does not exit normally. The memory management software detects some sort of problem. What's wrong and what is the error in the code? (Be specific and concise. You do not need to fix the problem – just explain it precisely.)

The memory manager reports a "double delete" error when the destructor for the second object is executed. The error is in the `Base::operator=` code. This assignment operator copies a pointer instead of creating a copy of the array. As a result, after the assignment `b1=d1`, both objects point to the single `ia_` array originally allocated to `d1`, and the destructors for `b1` and `d1` both attempt to delete it, causing the double delete error.

Question 4. (20 points) The always entertaining virtual function question. The
following program compiles, runs, and produces output with no error messages or other
problems. Answer questions about it on the next page.

```cpp
#include <iostream>
using namespace std;

class SuperThing {
public:
  virtual void m1() { m2(); cout << "super::m1" << endl; }
          void m2() {       cout << "super::m2" << endl; }
          void m3() {       cout << "super::m3" << endl; }
};
class Thing: public SuperThing {
public:
  virtual void m2() { m1(); cout << "thing::m2" << endl; }
};
class SubThing: public Thing {
public:
  virtual void m1() {       cout << "sub::m1" << endl; }
          void m3() { m2(); cout << "sub::m3" << endl; }
};

int main() {
  SuperThing *super = new Thing();
  Thing      *th    = (Thing*)super;
  SubThing   *sub   = new SubThing();
  Thing      *thsub = sub;

  ///// HERE /////

  cout << "---" << endl;
  th->m1();
  th->m3();
  cout << "---" << endl;
  sub->m1();
  sub->m3();
  cout << "---" << endl;
  thsub->m1();
  thsub->m3();

  return 0;
}
```

(Question continued on next page – you may remove this page if you wish.)

Question 4. (cont.) (a) (8 points) Complete the following diagram to show the runtime state of the program when execution reaches the comment `///// HERE /////` in function `main`. The diagram should include the variables in `main` (already supplied), the objects they point to, pointers from objects to their vtables, and pointers from vtables to the correct functions. To save time, boxes for the variables in main, the vtables, the functions, and the first object created by the program, have been provided for you. A couple of the arrows representing some of the pointers are also included to get you started. You need to supply all additional objects and pointers needed (if any). Be sure that the order of pointers in the various vtables is clear.
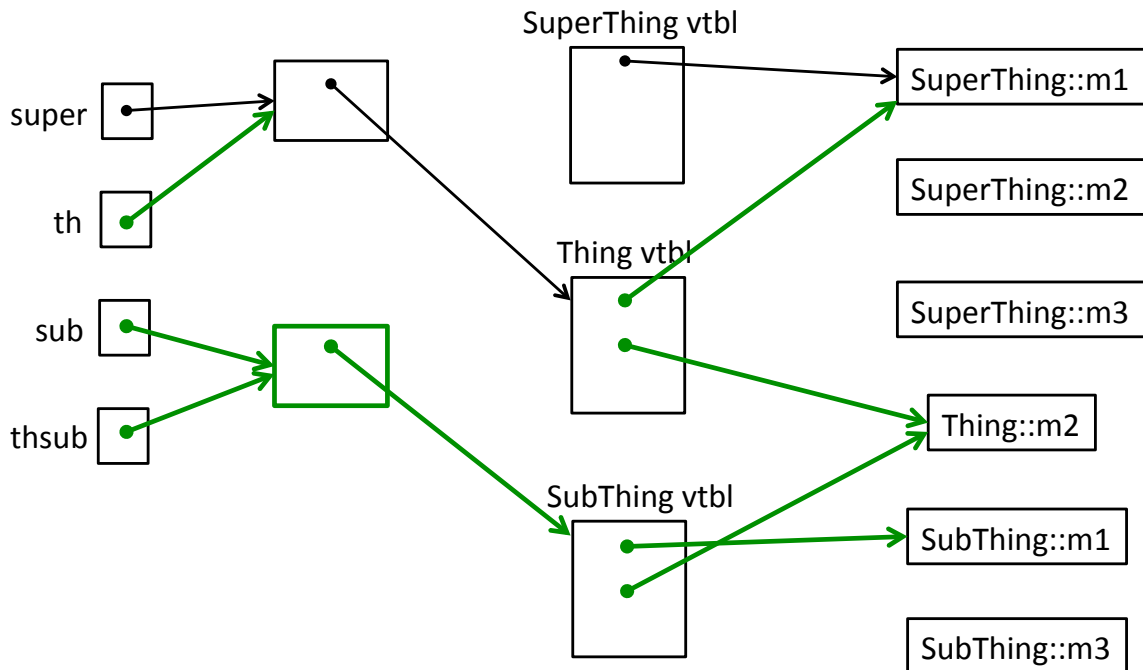
(b) (12 points) What does this program print when it is executed?

```
---
super::m2
super::m1
super::m3
---
sub::m1
sub::m1
thing::m2
sub::m3
---
sub::m1
super::m3
```

Question 5. (15 points) Concurrency. The arguments to the following program are a
string (`argv[1]`) and a list of one or more files (`argv[2]`, `argv[3]`, etc.). The
program does a parallel search to count the number of occurrences of the string in the
files. A separate worker thread is created to count each file and then update a global
total. Answer questions about this program following the end of the code on the next
page. **Do not** detach this page – you may want to write some answers on it.

```
#includes omitted - assume all necessary ones are provided

static char *search_word = NULL;   // Word to look for
static int nfound = 0;      // total number of occurrences of
                            // search_word in all files
static pthread_mutex_t sum_lock;  // added

// Return number of occurrences of search_word in file
// whose name is fname.  Does exit(1) on failure.
int numOccurences(char *fname) {
  /* Implementation omitted */
}
// Thread "starter" function: search file whose name is given in
// arg and add number of occurrences of search_word to nfound.
void *thread_main(void *arg) {
  char *fname = (char *)arg;
  int count = numOccurences(fname);
  printf("  File: %s - number found: %d\n", fname, count);
  pthread_mutex_lock(&sum_lock);    // added
  nfound += count;
  pthread_mutex_unlock(&sum_lock);  // added
  return NULL;
}
// Search for string (argv[1]) in files (argv[2], ...) and print
// total number of occurrences of that string in those files.
int main(int argc, char** argv) {
  if (argc < 2) {
    fprintf(stderr, "usage: %s STRING [FILE...]\n", argv[0]);
    return EXIT_FAILURE;
  }
  search_word = argv[1];
  int num_threads = argc - 2;
  printf("Searching for: %s\n", search_word);
  pthread_t *thds = malloc(sizeof(pthread_t) * num_threads);
  pthread_mutex_init(&sum_lock, NULL);   // added
  for (int i = 0; i < num_threads; i++) {
    if (pthread_create(&thds[i],NULL,&thread_main,argv[2+i])!=0){
      fprintf(stderr, "pthread_create failed\n");
      return EXIT_FAILURE;
    }
  }  // function main continued on next page
```

Question 5. (cont.) Listing of main program, continued.

```
  for (int i = 0; i < num_threads; i++) {
    if (pthread_join(thds[i], NULL) != 0) {
      fprintf(stderr, "pthread_join failed\n");
      return EXIT_FAILURE;
    }
  }
  printf("Total number found: %d\n", nfound);
  pthread_mutex_destroy(&sum_lock);     // added
  free(thds);
  return 0;
}
```

(a) (5 points) Are there any possible race conditions (possible synchronization errors) if this program is used to search a *single* file for the string? Give a precise explanation of why or why not.

No. Once `main` has created a second thread to search the file it waits for that thread to exit before continuing. So there is no possibility of concurrent access to any shared variables.

(b) (5 points) Are there any possible race conditions (possible synchronization errors) if this program is used to search *multiple* files for the string? (i.e., more than one file name is given in the program arguments) Give a precise explanation of why or why not.

Yes. All of the threads read global variable `nfound`, add something to the value read, and store it back. If the read-modify-write steps in two different threads are interleaved in the wrong order, one or more updates can be lost. (We are assuming that `numOccurrences` is thread-safe and that no two program arguments contain the same file name, which would introduce additional race conditions.)

(c) (5 points) If there are any race conditions (possible synchronization errors) in the code, modify or add missing pieces to the code to make it thread-safe. You should write your corrections on the code listing, showing what needs to be changed or inserted. You can declare additional global or local variables as needed. If no changes are needed, simply say so. Some possibly useful function prototypes:
- pthread_create(thread, attr, start_routine, arg)
- pthread_exit(status)
- pthread_join(thread, value_ptr)
- pthread_cancel (thread)
- pthread_mutex_init(pthread_mutex_t * mutex, attr) // attr=NULL usually
- pthread_mutex_lock(pthread_mutex_t * mutex)
- pthread_mutex_unlock(pthread_mutex_t * mutex)
- pthread_mutex_destroy(pthread_mutex_t * mutex)

Additional code needed to create and use a lock added in the code on the previous page. Look for the bold lines with the "`//added`" comments.

Question 6. (10 points) In the hw3 project we wrote a complex file structure to disk holding hash tables, indexes, and other information about collections of documents.

(a) (5 points) One of the very last things written to the file was a *checksum* field in the file header. What was the purpose of this field (briefly)?

It helps us verify that the file has not been damaged. The checksum is computed as a mathematical function of the contents of most of the file. If the checksum stored in the header doesn't match the value we get when we recalculate it when reading the file, we know the file has been corrupted.

(b) (5 points). The header of the file also contained a "magic number" (in hex, 0xCAFEF00D) and we were very careful to write this at the very end, after everything else had been written. Why? Why wait until the end instead of writing it earlier?

This is written last so that if our program crashes before we have finished writing the file the magic number won't be present. When a program reads the file, if the magic number is missing we know the file is incomplete or damaged.

Question 7. (6 points) Our favorite summer intern, A. Hacker, has been messing around in the network code and has decided that the TCP/IP networking software layers have too much overhead. Since TCP chops messages up into small packets to be delivered by IP, he proposes to eliminate TCP sockets from the code and simply send long messages as multiple packets directly, thereby saving the cost of using TCP. Is this a reasonable change? Would we be missing anything if we simply chopped the messages into packets and sent them via IP? If so, what would be missing? (Keep it brief and describe the main issue(s).)

Two main things that TCP does that would be missing if we sent packets directly:
- **Check for lost packets and arrange to have them retransmitted.**
- **When packets are received out of order, arrange them in the correct order to reassemble the original message properly.**

Notes: Remember that IP does not guarantee that packets will be received in the order they were transmitted, and packet delivery is on a "best efforts" basis, so some packets may be lost in transmission.

In addition to correctness, TCP also manages flow control to match the speed at which packets are sent to network capacity, and it also provides port numbers and some other features. But these are secondary compared to providing reliable, in-order transmission of the packets that comprise a message.

Question 8. (1 free point) Which of the following best matches your opinion about daylight savings time – the reason all the clocks moved forward one hour on Sunday morning. (Circle)

(a) It's great! There's still sunlight when I emerge from the basement to go home.

(b) Really didn't like having to get up an hour earlier for my 8:30 exam this Monday.

(c) Daylight? What is this "daylight" of which you speak?

(d) Love it. Means spring break can't be too far behind.

(e) Stupid idea. Let's leave all the clocks set to standard time.

(f) Everyone should use a sundial. Noon is when the sun is directly overhead.

(g) Stupid idea. All clocks should be set to GMT universal time so we don't have to deal with time zones.

(h) Time zones? Isn't that something that went out with the end of passenger railroads?

(i) Meh. Might get used to it by late summer, and then they just change it back again.

(j) Good reminder that it's time to get the yacht out of storage and go boating again.

(k) Don't care. I'm going to catch up on sleep during spring break so it won't matter what it says on the clock.

(l) Great idea. I'd never remember to replace the smoke detector batteries without it.

(m) It's great! But *not* the day before finals week!!

(x) I don't have an opinion, but please give me my free point anyway.

(y) I do have an opinion, but I'm not going to tell. Please give me my free point anyway.

(z) None of the above, but I still want my free point. My real opinion about daylight savings time is:

All answers received 1 point.

Have a great time over spring break!! See you next quarter!!!
The CSE 333 staff