

CSE 333 FINAL

Last Name:

First Name:

Student ID Number:

Name of person to your Left | Right

All work is my own. I had no prior knowledge of the exam content, nor will I share the content with others in CSE 333 who haven't taken it yet. Violation of these terms could result in a failing grade. (please sign)

Last Name:	
First Name:	
Student ID Number:	
Name of person to your Left	Name of person to your Right
All work is my own. I had no prior knowledge of the exam content, nor will I share the content with others in CSE 333 who haven't taken it yet. Violation of these terms could result in a failing grade. (please sign)	

Do not turn the page until you are told to do so.

Instructions

- This exam contains 14 pages, including this cover page. Show scratch work for partial credit, but put your final answers in the boxes and blanks provided.
- The last page is a reference sheet. *Detach it from the rest of the exam when you are told to do so.*
- The exam is closed book (no laptops, tablets, wearable devices, or calculators). You are allowed one page (US letter, double-sided) of *handwritten* notes.
- Please silence and put away all cell phones and other mobile or noise-making devices. Remove all hats, headphones, and watches.
- You have 110 minutes to complete this exam.

Advice

- Read questions carefully before starting. Skip questions that are taking a long time.
- Read *all* questions first and start where you feel the most confident.
- Relax. You are here to learn.

Question	1	2	3	4	5	6	Total
Possible Points	12	18	20	21	20	20	111

Question 1: Potpourri – Nice to Smell, Hard to Spell [12 pts]

- (A) Assume D is a derived class of class B. If we have `B* ptr = new B;`. If we try to assign `D* new_ptr = ?_cast<D*>(ptr);` will the following C++ casts compile ("C") or cause a compiler error ("E")? [4 pt]

static_cast _____ dynamic_cast _____
const_cast _____ reinterpret_cast _____

- (B) For the following scenarios, indicate whether it will result in a double delete ("D"), memory leak ("M"), or no issue ("N"). [4 pt]

<code>unique_ptr<int> p1(new int(1)); p1.release();</code>	_____
<code>unique_ptr<int> p2(new int(2)); p2.reset();</code>	_____
<code>shared_ptr<int> p3(new int(3)); shared_ptr<int> p4(p3.get());</code>	_____
<code>shared_ptr<int> p5(new int(5)); weak_ptr<int> p6(p5); p6.lock();</code>	_____

- (C) Assume we have the following ThreadMain function that gets run by 2 threads. Which of the following final values of g are possible ("Y"/"N") after joining all threads? [4 pt]

```
int g = 3;  
void* ThreadMain(void* args) {  
    g = g + 3;  
    g = g - 3;  
    return nullptr;  
}
```

-3 _____ 0 _____ 6 _____ 9 _____

Question 2: All in a Day's NET-WORK [18 pts]

- (A) What does a **positive (>0) return value** mean for the following network programming functions? Where applicable, you can write "not possible." [8 pt]

socket() :
connect() :
accept() :
write() :

- (B) Fill in the blanks using the word bank below: [6 pt]

addrinfo	AF_INET	AF_INET6	
in_addr	in6_addr	in_port_t	sa_family_t
sockaddr	sockaddr_in	sockaddr_in6	sockaddr_storage

An **IPv4** address (*just the address*) is stored in a struct _____.

What allows us to differentiate an IPv4 vs. IPv6 address? _____.

DNS results are found in a linked list of struct _____.

- (C) Name one piece of metadata that might be included in the header of the packet portion associated with the following network layers. [4 pt]

Network layer:
Transport layer:

Question 3: The KEY to Student Success [20 pts]

Student data uses a variety of identifiers, and systems often don't agree on which to use. For example, Canvas uses both UW **netid** (e.g., hhusky) and a special **Canvas ID** (e.g., 4012345), while Gradescope uses **email** (e.g., hhusky@uw.edu). Both systems use comma-separated values files (.csv), but we need them to talk to each other! We will use the following files (*i.e.*, what the columns in each row are):

- Canvas roster file: "netid,canvas_id,name,section"
- Gradescope grade file: "first_name,last_name,SID,email,section,score,..."
- Canvas upload file: "canvas_id,score,comment"

We will use **C++ STL** to combine a Canvas roster file and a Gradescope grade file to produce a corresponding Canvas upload file. The top of our file is shown below:

```
#include <iostream>
#include <fstream>
#include <map>
#include <string>
#include <boost/algorithm/string.hpp>
using namespace std;
```

- (A) Complete the helper function below; *you may not need to fill in all blanks*. `in.getline()` returns false on an error. The first two arguments to `boost::split` are `vector<string>&` output and `const string&` input. [4 pt]

```
static const int kBufSize = 1024;
char buf[kBufSize]; // buffer for getline()

// Reads line of CSV and splits by commas into output parameter data.
// Returns true if successful and false on bad read.
bool GetRow(ifstream& in, vector<string>* data) {
    if (in.getline(buf, kBufSize)) {
        _____;

        boost::split(_____,
                    _____,
                    boost::is_any_of(","),
                    boost::token_compress_off);

        _____;
    } else {
        _____;
    }
}
```

(B) Complete the following snippet of `main` to read the Canvas roster file `r_file` and generate a mapping from `netid` to Canvas ID. [6 pt]

- You may assume that all rows in the file are properly formatted and have unique IDs.

```
ifstream r_file(argv[1], ifstream::in); // open Canvas roster file
_____ canvas_ids; // map netids to canvas ids
_____ row_data;

while (GetRow(_____, _____)) {
    // Assumes row is "netid,canvas_id,name,section"
    _____;
}
```

(C) Complete the following snippet of `main` that comes after the code in Part B to read the Gradescope grades file `g_file` and output the corresponding Canvas upload file to `stdout`. In addition to the provided blanks, the body of the while-loop is left for you to fill in. [10 pt]

- You may assume that all rows in the file are properly formatted; however, you may NOT assume that all students in the grades file are found in the Canvas roster (*i.e.*, students drop).
- You can access `kBufSize`, `buf` (global), `canvas_ids`, and `row_data` (Part B).
- Useful `string` member functions to extract a `netid` from email:
 - `string substr(size_t pos, size_t len)`; returns the substring of length `len` starting from position `pos`.
 - `size_t find(char* s)`; returns the position of the first occurrence of `s`.

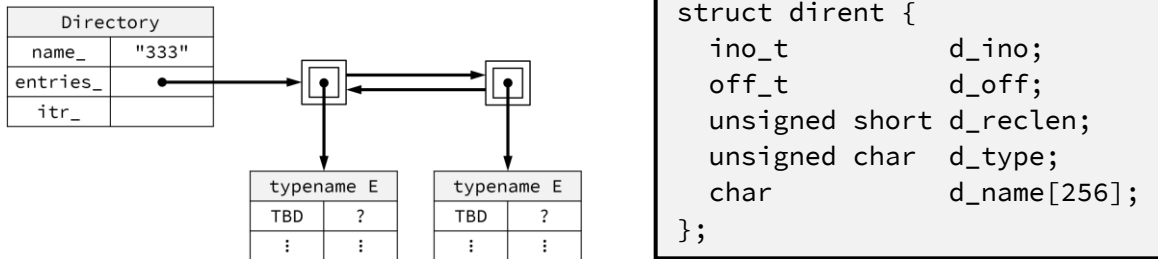
```
ifstream g_file(argv[2], ifstream::in); // open Gradescope file
_____; // discard header row

while (GetRow(_____, _____)) {
    // Assumes row is "first_name,last_name,SID,email,section,score,..."
    // Outputs rows "canvas_id,score,comment" (leave comment blank)
}

}
```

Question 4: SMART Revisions [21 pts]

We want to create a C++ wrapper class for directories to abstract away the POSIX library details. For all parts of this question, you may assume the appropriate C++ and POSIX libraries are included, along with `using namespace std;`.



As seen above, a `Directory` object should have private fields for (1) the name of the directory, (2) a doubly-linked list of directory entries, and (3) an iterator to the “current” entry within the directory. We want to be able to use a wrapper class for directory entries, so we will make this a template class. To account for the unknown directory entry size, we will use an STL container of smart pointers (`unique_ptr`) to reduce the cost of container copying and automatically handle memory management.

(A) Fill in the data member types below for `Directory` to match the above description. [3 pt]

```
template <typename E> class Directory {
public:
    // declare/define member functions

private:
    _____ name_;
    _____ entries_;
    _____ itr_;
}; // class Directory
```

(B) Give an inline definition (*i.e.*, within the class declaration in a `.h` file) for `get_size()`, which returns the number of entries in the directory. [3 pt]

(C) Complete the definition of the 1-arg `Directory` constructor on the opposite page, which will open the directory and build the entries list from the directory entries. [5 pt]

- `typename E` is assumed to have a 1-arg constructor that takes a `dirent`.
- The iterator should be initialized to the first entry of the directory.
- Most of the POSIX code is provided.
- You must include an initializer list.

```

template <typename E>
Directory::Directory(string name) : _____ {
    DIR* dirp = opendir(name.c_str());
    if (!dirp) {
        cerr << "Error opening directory" << name << endl;
        exit(EXIT_FAILURE);
    }
    dirent* entry;
    while ( (entry = readdir(dirp)) ) {
        _____;
    }
    _____;
    _____;
}

```

- (D) Complete the definition of `ReadDir` below, which mimics the `readdir()` POSIX function. To preserve the privacy of entries, we return a raw pointer to a heap-allocated copy. You may assume that `E` has a default constructor. [6 pt]

```

// Returns ptr to heap copy of current dir entry (nullptr if at
// end), then moves to next dir entry, if possible.

template <typename E> _____ Directory::ReadDir() {
    if ( _____ ) {
        return nullptr; // at end of directory
    } else {
        _____; // heap allocate
        _____; // copy
        _____; // advance
        _____; // return
    }
}

```

- (E) Give two lines of C++ code to (1) create a local `Directory` object for directory "333", using the POSIX `dirent` for directory entries, and (2) call `ReadDir` without leaking memory. [4 pt]

Question 5: Linguistic INHERITANCE [20 pts]

Consider the following C++ **classes**, which compile without errors.

```
class B {
public:
    void f1() { cout << "B::f1, "; }
    virtual void f2() { f1(); cout << "B::f2, "; }

private:
    int b_ = 1969;
};

class C : public B {
public:
    virtual void f1() { cout << "C::f1, "; }
    virtual void f3() { cout << "C::f3, "; }

private:
    int c_ = 1972, ansi_ = 1989;
};


class CPP : public C {
public:
    void f2() { f3(); cout << "CPP::f2, "; }

private:
    int* cpp_ = new int(1985);
};
```

(A) Draw a conceptual/memory diagram of a default-constructed object of class C below. Don't show vptr's. [2 pt]

(B) The code `B* ptr = new CPP; delete ptr;` causes a memory leak! *Briefly* describe what you would add to the class definitions to fix this issue. [4 pt]

(C) Complete the **virtual function table diagram** below by adding the remaining class methods (whether dispatched statically or dynamically) on the right and then drawing the appropriate function pointers from the vtables. *Ordering of the function pointers matters!* One is already included for you. [6 pt]



(D) Assume we have objects and pointers as defined below. For each row of the table, **fill in the result** on the right with the corresponding `stdout` output, “compiler error,” or “runtime error.” [8 pt]

```
B* bc = new C;
B* bp = new CPP;
C* cp = new CPP;
```

<code>bc->f1();</code>	
<code>bc->f2();</code>	
<code>bc->f3();</code>	
<code>bp->f2();</code>	
<code>cp->f1();</code>	
<code>cp->f2();</code>	

Question 6: Results Hanging by a THREAD [20 pts]

We are writing a *multi-threaded C++ program* that will total the ballots in an election between Candidates A and B (stored in an `int` array). The top of the file is given below:

```
#include <pthread.h>
#include <iostream>

static const int kNumThreads = 2;
static const int kNumBallots = 100;

// global ballot count indices: 0 = no vote, 1 = A, 2 = B
static int counts[] = {0, 0, 0};
static pthread_mutex_t counts_lock;

struct thd_arg {
    int* ptr; // ballot array pointer
    int num; // number of ballots to count
};
```

- (A) Complete the snippet of `main` below that creates the children threads and splits the ballot counting work evenly between them. The function that the children threads will run is called `ThreadMain`. For simplicity, skip error checking and assume `kNumBallots` is a multiple of `kNumThreads`. The passed argument should be *heap-allocated*. [6 pt]

```
int ballots[] = {...}; // ballots array
pthread_t thds[kNumThreads]; // array of thread ids
pthread_mutex_init(&counts_lock, nullptr);

for (int i = 0; i < kNumThreads; ++i) {

    thd_arg* args = _____;

    args->num = _____;

    args->ptr = _____;

    pthread_create(_____,
                  _____);

}
```

- (B) Complete the thread safe function `ThreadMain` on the opposite page. All elements in the ballot array contain the values 0, 1, or 2, which should be tallied in their corresponding indices in `counts[]`. Assume `arg` is heap-allocated. [6 pt]

```

// Reads num ballots from the ballot array and adds the votes
// to the corresponding indices of the counts[] global array.
// Takes ownership of arg. Assumes counts_lock is initialized.
void* ThreadMain(void* arg) {
    thd_arg* a = static_cast<thd_arg*>(arg);

    for (_____ ) {
        _____;
        _____;
        _____;
    }

    _____;

    return _____;
}

```

- (C) Assume `kNumBallots` is 100, and we have two CPUs available to us. Will the multithreaded code perform better when `kNumThreads` is 2 or 100? *Briefly* explain. [4 pt]

- (D) It turns out that we can improve our multithreaded performance by changing the way we lock in the code! *Briefly* describe how we accomplish this and why it helps our performance. Hint: It will require 3 mutexes. [4 pt]

Change:
Why It Helps:

**THIS PAGE PURPOSELY
LEFT BLANK**

CSE 333 Reference Sheet (Final)

C Library Header – stdlib.h

```
EXIT_SUCCESS // success termination code
EXIT_FAILURE // failure termination code

void exit (int status); // terminate calling process
```

Error Library – errno.h

```
errno // # of the last error, usually checked against defined consts

EAGAIN // try again
EBADF // bad file/directory/socket descriptor
EINTR // interrupted function
EWOULDBLOCK // operation would block
```

C++ Standard Template Library – vector, list, map, etc.

```
.begin() // get iterator to beginning (first element)
.end() // get iterator to end (one past last element)
.size() // get container size
.erase(...) // erase element (pass 1 iterator) or range (pass 2 iterators)

template <class T> class std::vector;
    • .operator[](), .push_back(), .pop_back()
template <class T> class std::list;
    • .push_back(), .pop_back(), .push_front(), .pop_front(), .sort()
template <class Key, class T> class std::map;
    • .operator[](), .insert(), .find(), .count()
template <class T1, class T2> struct std::pair
    • .first, .second
```

C++ STL Algorithms – algorithm

```
std::find() // returns iterator to element in range that matches val
std::for_each() // apply function to each element in range
std::min_element() // get iterator to smallest element in range
std::max_element() // get iterator to largest element in range
std::sort() // sorts range into ascending order
```

C++ Smart Pointers Library – memory

```
template <class T> class unique_ptr;
    • .get(), .reset(), .release()
template <class T> class shared_ptr;
    • .get(), .use_count(), .unique()
template <class T> class weak_ptr;
    • .lock(), .use_count(), .expired()
```

POSIX Headers – unistd.h, arpa/inet.h, netdb.h

```
ssize_t read (int fd, void* buf, size_t count);
ssize_t write (int fd, const void* buf, size_t count);
int getaddrinfo (const char* hostname, const char* service,
                 const struct addrinfo* hints, struct addrinfo** res);
int socket (int domain, int type, int protocol);
int connect (int fd, const struct sockaddr* addr, socklen_t addrlen);
int bind (int sockfd, const struct sockaddr* addr, socklen_t addrlen);
int listen (int sockfd, int backlog);
int accept (int sockfd, struct sockaddr* addr, socklen_t* addrlen);
```

Pthreads Header – pthread.h

```
pthread_t          // data type to identify a thread
pthread_mutex_t    // data type for a mutex

int pthread_create (pthread_t* thread, const pthread_attr_t* attr,
                   void* (*start_routine)(void*), void* arg);
int pthread_join (pthread_t thread, void** retval);
int pthread_detach (pthread_t thread);

int pthread_mutex_init (pthread_mutex_t* mutex,
                       const pthread_mutexattr_t* attr);
int pthread_mutex_lock (pthread_mutex_t* mutex);
int pthread_mutex_unlock (pthread_mutex_t* mutex);
int pthread_mutex_destroy (pthread_mutex_t* mutex);
```