

Lecture 17

RAII & Memory Management

CS211 – Fundamentals of Computer Programming II
Branden Ghen a – Fall 2021

Slides adapted from:

Jesse Tov (Northwestern), Hal Perkins (Washington), Godmar Back (Virginia Tech)

Administrivia

- Project specifications due today
 - Each group gets assigned a shepherd
 - Shepherd will provide feedback on your specification items this week
- In the meantime: get started on code now!
 - Get little bits of functionality working this week
 - Make sure you've made significant progress *before* Thanksgiving break

Lecture plan from here

- Thursday 11/18
 - More GE211 examples: finish up before + animation
- Tuesday 11/23
 - Bonus lecture: Version control and Git
- Tuesday 11/30
 - C and C++ wrapup
- Thursday 12/2
 - Final project demos for anyone interested

Today's Goals

- Consider the RAII programming idiom:
(Resource Acquisition Is Initialization)
 - Understand how it is making development easier in C++
- Discuss C++ memory management
 - What exists and how it works
 - How to use Smart Pointers to make it easy too

Outline

- **C++ Strings**
- RAII
- C++ Memory Management
- Smart Pointers

Strings in C++

- Everything you wanted from C strings and didn't get

```
#include <string>
```

```
std::string s1 = "Test";
```

```
s1 += " String";
```

```
s1[0] = 'B';
```

```
std::cout << s1 << "\n"; // prints "Best String"
```

C++ string operations

- Iterators
 - Including reverse and constant
- Sizing
 - Characters and memory
- Access to characters

Iterators:

begin	Return iterator to beginning (public member function)
end	Return iterator to end (public member function)
rbegin	Return reverse iterator to reverse beginning (public member function)
rend	Return reverse iterator to reverse end (public member function)
cbegin <small>C++11</small>	Return const_iterator to beginning (public member function)
cend <small>C++11</small>	Return const_iterator to end (public member function)
crbegin <small>C++11</small>	Return const_reverse_iterator to reverse beginning (public member function)
crend <small>C++11</small>	Return const_reverse_iterator to reverse end (public member function)

Capacity:

size	Return length of string (public member function)
length	Return length of string (public member function)
max_size	Return maximum size of string (public member function)
resize	Resize string (public member function)
capacity	Return size of allocated storage (public member function)
reserve	Request a change in capacity (public member function)
clear	Clear string (public member function)
empty	Test if string is empty (public member function)
shrink_to_fit <small>C++11</small>	Shrink to fit (public member function)

Element access:

operator[]	Get character of string (public member function)
at	Get character in string (public member function)
back <small>C++11</small>	Access last character (public member function)
front <small>C++11</small>	Access first character (public member function)

C++ string operations

- Modification of strings
 - Add or remove from them
- Operations
 - Get C string from `std::string`
 - Find
 - Substring
 - Compare

Modifiers:

<code>operator+=</code>	Append to string (public member function)
<code>append</code>	Append to string (public member function)
<code>push_back</code>	Append character to string (public member function)
<code>assign</code>	Assign content to string (public member function)
<code>insert</code>	Insert into string (public member function)
<code>erase</code>	Erase characters from string (public member function)
<code>replace</code>	Replace portion of string (public member function)
<code>swap</code>	Swap string values (public member function)
<code>pop_back</code> <small>C++11</small>	Delete last character (public member function)

String operations:

<code>c_str</code>	Get C string equivalent (public member function)
<code>data</code>	Get string data (public member function)
<code>get_allocator</code>	Get allocator (public member function)
<code>copy</code>	Copy sequence of characters from string (public member function)
<code>find</code>	Find content in string (public member function)
<code>rfind</code>	Find last occurrence of content in string (public member function)
<code>find_first_of</code>	Find character in string (public member function)
<code>find_last_of</code>	Find character in string from the end (public member function)
<code>find_first_not_of</code>	Find absence of character in string (public member function)
<code>find_last_not_of</code>	Find non-matching character in string from the end (public member function)
<code>substr</code>	Generate substring (public member function)
<code>compare</code>	Compare strings (public member function)

Strings with different character sizes

- All are actually implementations of the generic `std::basic_string`
 - 16-bit “wide” characters
 - Strings of 8-bit, 16-bit, or 32-bit characters

Several typedefs for common character types are provided:

Defined in header `<string>`

Type	Definition
<code>std::string</code>	<code>std::basic_string<char></code>
<code>std::wstring</code>	<code>std::basic_string<wchar_t></code>
<code>std::u8string</code> (C++20)	<code>std::basic_string<char8_t></code>
<code>std::u16string</code> (C++11)	<code>std::basic_string<char16_t></code>
<code>std::u32string</code> (C++11)	<code>std::basic_string<char32_t></code>

- UTF-8 mostly works with `std::string` by default
 - Some helper functions won't work properly though...
 - Needs additional libraries for many functions

Outline

- C++ Strings
- **RAII**
- C++ Memory Management
- Smart Pointers

RAII-structured libraries enable simple dynamic memory

- `std::vector`, `std::string`, and other library containers must use dynamic memory internally
 - But we never have to call `vector.destroy()` or `free(string)`
- What makes memory management so automatic in C++?
- Programming paradigm: RAII
 - Resource Acquisition Is Initialization
 - Basic idea:
 - Wrap resources in an object
 - Allocate when you initialize and deallocate when destructed

What is a “resource”?

- Abstractly:
 - Something you need to get your computation done,
 - That you can run out of,
 - So you need to keep track of what you’re using and release what you aren’t
- Concretely:
 - Memory!
 - File handles
 - Network sockets
 - Database sessions
 - Acquired *locks* (concurrent programming)

The problem: leaking resources

```
#include <cstdio >
```

```
void handle_file(std::string const& name) {  
    FILE *f = fopen(name.c_str(), "r");  
  
    // various code here using the file  
  
}
```

Didn't close the file!
There's a resource leak!!

The problem: leaking resources

```
#include <cstdio >

void handle_file(std::string const& name) {
    FILE *f = fopen(name.c_str(), "r");

    // various code here using the file

    if (some error occurred) { return; }

    // various more code using the file

    fclose(f);
}
```

What's wrong here?

The problem: leaking resources

```
#include <cstdio >
```

```
void handle_file(std::string const& name) {  
    FILE *f = fopen(name.c_str(), "r");  
  
    // various code here using the file  
  
    if (some error occurred) { return; }  
  
    // various more code using the file  
  
    fclose(f);  
}
```

More common cause: early returns
Always beware when code returns early

Exceptions make early returns even worse

```
void helper() {  
    if (some problem detected) { throw std::runtime_error("Oops"); }  
  
    // various code here using the file  
}
```

```
void handle_file(std::string const& name) {  
    FILE *f = fopen(name.c_str(), "r");  
  
    // various code here using the file  
  
    helper(); // might throw an exception never "return"  
  
    // various more code using the file  
  
    fclose(f);  
}
```

Can't clean up here without try/catch everywhere

C++ solution: Resource Acquisition Is Initialization

- Never open/close or free/allocate manually
- Instead make a class
 - Allocate in the constructor
 - Programmer calls this when initializing the object variable
 - Deallocate in the **destructor**
 - Automatically occurs. Programmer doesn't have to do anything!

Destructors

- Same concept as constructors: used to clean up an object
 - Automatically called when the object goes out of scope
 - Note: you never call the destructor yourself!
- Handles any cleanup, including freeing necessary resources

```
std::ifstream::~~ifstream() {  
    // close the file here  
}
```

Destructors allow resources to automatically be cleaned up

```
#include <fstream>
```

```
void handle_file(std::string const& name) {  
    std::ifstream f(name, "r");  
  
    // do stuff with the file  
  
} // f.~ifstream() happens automatically here
```

Destructors allow resources to automatically be cleaned up

```
#include <fstream>

void handle_file(std::string const& name) {
    std::ifstream f(name, "r");

    // do stuff with the file.
    // Possibly return or throw exceptions!
} // f.~ifstream() happens here regardless
```

The destructor is guaranteed to run.
Even if there is an exception!

Break + What might the `std::string` implementation look like?

```
class Owned_string {  
public:  
    Owned_string (); // empty string  
    Owned_string(const char* cstr); // allocates memory  
    ~Owned_string (); // frees memory  
  
private:  
    std:: size_t size_; // logical size of string  
    std:: size_t capacity_; // allocated size of `data_`  
    char *data_; // ptr to char array (or null)  
};
```

Outline

- C++ Strings
- RAII
- **C++ Memory Management**
- Smart Pointers

C++ memory management

- In C, dynamic memory was very important for making any realistic program that responds to user input
- In C++, because of RAII concepts and the Standard Template Library, we haven't had to manually use dynamic memory at all!
 - But it is still there, happening
 - And we could harness it ourselves if we need to

Reminder: C memory allocation

```
void* malloc(size_t size)
```

- Requests `size` bytes of memory from the heap
- Returns a pointer to this new **object**
 - Not associated with any variable (sort of like string literals)
 - It has no value by default
- The object persists until it is manually deallocated
 - Deallocated through a call to `free()`

C++ memory allocation

- Allocate with the `new` keyword and a type
 - No need to specify number of bytes anymore
 - Works for primitive types and for objects
 - Examples:
 - `int* value_ptr = new int;`
 - `Posn<int>* p = new Posn<int>;`
- Deallocate with the `delete` keyword and the pointer
 - Example: `delete p;`
- Warning: never mix-and-match `malloc()/free()` with `new/delete`
 - **UNDEFINED BEHAVIOR** (`free()` doesn't call destructor!!)

Dynamic arrays in C++

- For `new`, add the size of the array after the type

```
int* data = new int[10];
```

- For `delete`, must instead use `delete []`

- **Important:** Must remember this or **UNDEFINED BEHAVIOR** 🤖

- Reason:

- `delete` calls the destructor and then frees the memory
- `delete []` iteratively calls destructors and then frees the memory

- `delete []` *could* have worked for everything, but it would be less efficient

Dynamic arrays in C++

- For new, add the size of the array after the type

```
int arr = new int[10];
```

Just use `std::array` or `std::vector` instead!!

- Reason:

- `delete` calls the destructor and then frees the memory
- `delete[]` iteratively calls destructors and then frees the memory

- `delete[]` *could* have worked for everything, but it would be less efficient

C dynamic memory vs C++ dynamic memory

	<code>malloc()</code>	<code>new</code>
What is it?	a function	an operator or keyword
How often used (in C)?	often	never
How often used (in C++)?	rarely	sometimes (often, but by a library without the dev knowing)
Allocated memory for	anything	arrays, structs, objects, primitives
Returns	a <code>void*</code> <i>(should be cast)</i>	appropriate pointer type <i>(doesn't need a cast)</i>
When out of memory	returns <code>NULL</code>	throws an exception
Deallocating	<code>free()</code>	<code>delete</code> or <code>delete[]</code>

Null pointers in C

- While NULL still works (legacy from C), there's a better way
- `nullptr` is the preferred literal
 - Same meaning as NULL, but its type is explicitly `T*` for *any* type `T`
 - Still converts to 0 when needed

- C++ example:

```
void print(int* value_ptr);  
void print(int value);
```

```
print(NULL); // calls print for type int 🤖  
print(nullptr); // calls print for type int* 😎
```

Outline

- C++ Strings
- RAII
- C++ Memory Management
- **Smart Pointers**

Using dynamic memory in a class

- Constructor will call `new` to allocate memory for some data member
- Destructor will call `delete` to free the memory when the object goes out of scope
- Observation:
 - Memory is manually created and initialized to values
 - But deletion is almost always just calling `delete`
 - We could use RAII to do this for us

C++ Smart Pointers

- A smart pointer is an object that stores a pointer to a heap-allocated object
 - Behaves just like a normal C++ pointer by overloading `*`, `->`, `[]`, etc.
- Smart pointers do the memory management for you
 - Automatically deletes the pointed-to object if the smart pointer goes out of scope
 - I.e., if the memory would leak, it is instead freed
- Smart pointers are the modern C++ way to do dynamic memory

Unique pointer (unique_ptr)

- Takes ownership of a pointer
- Allows access to the value pointed to
- Invokes `delete` automatically
 - Either when the `unique_ptr` goes out of scope via the destructor
 - Or when the owned pointer is overwritten

```
#include <memory>
```

```
std::unique_ptr<char> letter_ptr(new char('a'));
```

```
char letter = *letter_ptr; // sets letter to 'a'
```

Smart pointers are automatically freed

```
#include <memory>

void handle_memory() {
    std::unique_ptr<double> d(new double(3.7));

    // do stuff with the pointer
    // Possibly return or throw exceptions!
} // memory is freed here regardless
```

The destructor is guaranteed to run.
Even if there is an exception!

Unique_ptr ownership rules

- Matches the ownership rules we discussed previously
 - There is only one single owner of a unique_ptr
 - Which in turn owns the memory

- Cannot be copied

```
std::unique_ptr<int> x(new int(5)); // OK
std::unique_ptr<int> y(x); // Fails, no copy constructor
std::unique_ptr<int> z; // OK, holds nullptr
z = x; // Fails, no assignment operator
```

- Ownership can be transferred if needed
 - release() gives up ownership of the pointer
 - reset() deletes the current pointer (if any) and stores a new one

Unique_ptr and arrays

- unique_ptr can store arrays as well
 - Will call delete[] on destruction

```
int main() {  
    std::unique_ptr<int[]> x(new int[5]);  
    x[0] = 1;  
    x[1] = 2;  
  
    return 0; // memory will be freed automatically  
}
```

Shared pointers (`shared_ptr`)

- Similar to a `unique_ptr`, except that there can be multiple owners
 - Different ownership policy
- Tracks the number of owners to decide when to free
 - Copy/assign operators do work and increment number of owners
 - Destructor decrements number of owners
 - Frees memory if number of owners hits zero
- Technique is known as “reference counting”
 - Higher overhead than a `unique_ptr` has: slower to use

Main takeaways

- Smart pointers are how memory is managed in modern C++
 - Still have to use `new` operator, but never need to `delete`
- `unique_ptr` automatically manages ownership rules for us
 - Ensures that there is only one owner at a time
 - Ensure that memory is properly freed if there would be no owner

Outline

- C++ Strings
- RAII
- C++ Memory Management
- Smart Pointers