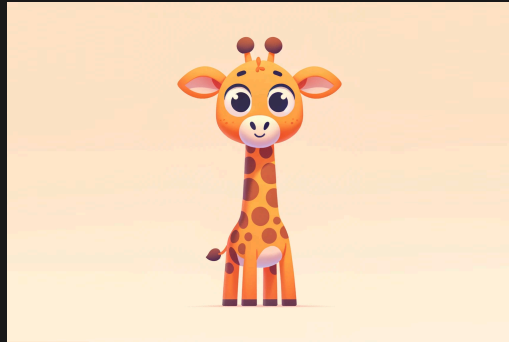


COMP6771

Smart Pointers

Lecture 5.3

Author(s): Hayden Smith



[\(Download as PDF\)](#)

In This Lecture

- **Why?** 🤔
 - Managing unnamed / heap memory can be dangerous, as there is always the chance that the resource is not released / free'd properly. We need solutions to help with this.
- **What?** 📄
 - Smart pointers
 - Unique pointer, shared pointer
 - Partial construction



Object Lifetimes

To create safe object lifetimes in C++, we always attach the lifetime of one object to that of something else

- Named objects:
 - A variable in a function is tied to its scope
 - A data member is tied to the lifetime of the class instance
 - An element in a `std::vector` is tied to the lifetime of the vector
- Unnamed objects:
 - A heap object should be tied to the lifetime of whatever object created it
 - Examples of bad programming practice
 - An owning raw pointer is tied to nothing
 - A C-style array is tied to nothing



RAII

RAII = Resource Acquisition Is Initialization

In summary, resource management was really about emphasising RAII

- Resource = heap object
- A concept where we encapsulate resources inside objects
 - Acquire the resource in the constructor
 - Release the resource in the destructor
 - eg. Memory, locks, files
- Every resource should be owned by either:
 - Another resource (eg. smart pointer, data member)
 - Named resource on the stack
 - A nameless temporary variable



Making A Pointer Safe

We could write a class to make a pointer safe.

```
1 // myintptr.h
2
3 class MyIntPtr {
4 public:
5     // This is the constructor
6     MyIntPtr(int* value): value_{value} {}
7
8     // This is the destructor
9     ~MyIntPtr() {
10         // Similar to C's free function.
11         delete value_;
12     }
13
14     int* GetValue() {
15         return value_
16     }
17
18 private:
19     int* value_;
20 };
```

```
1 void fn() {
2     // Similar to C's malloc
3     MyIntPtr p{new int{5}};
4     // Copy the pointer;
5     MyIntPtr q{p.GetValue()};
6     // p and q are both now destructed.
7     // What happens?
8 }
```



Smart Pointers

- Smart pointers are ways of wrapping unnamed (i.e. raw pointer) heap objects in named stack objects so that object lifetimes can be managed much more safely
- Introduced in C++11
- Usually two ways of solving problems
 - `unique_ptr` + raw pointers
 - `shared_ptr` + `weak_ptr`

Type	Shared ownership	Take ownership
<code>std::unique_ptr<T></code>	No	Yes
raw pointers	No	No
<code>std::shared_ptr<T></code>	Yes	Yes
<code>std::weak_ptr<T></code>	No	No



Unique Pointer

- `std::unique_ptr<T>`
 - The unique pointer owns the object
 - When the unique pointer is destructed, the underlying object is too
- raw pointer (observer)
 - Unique pointer may have many observers
 - There is an appropriate use of raw pointers (or references) in C++
 - Once the original pointer is destructed, you must ensure you don't access the raw pointers (no checks exist)
 - Those observers do not have ownership over the pointer

Also note the use of `nullptr` in C++ instead of `NULL`.

Unique Pointer

```
1 #include <memory>
2 #include <iostream>
3
4 int main() {
5     auto up1 = std::unique_ptr<int>{new int};
6     // auto up2 = up1; // no copy constructor
7     std::unique_ptr<int> up3;
8     // up3 = up2; // no copy assignment
9
10    up3.reset(up1.release()); // OK
11    auto up4 = std::move(up3); // OK
12    std::cout << up4.get() << "\n";
13    std::cout << *up4 << "\n";
14    std::cout << *up1 << "\n";
15 }
```

unique.cpp

Observer Pointer

```
1 #include <memory>
2 #include <iostream>
3
4 int main() {
5     auto up1 = std::unique_ptr<int>{new int{0}};
6     *up1 = 5;
7     std::cout << *up1 << "\n";
8     auto op1 = up1.get();
9     *op1 = 6;
10    std::cout << *op1 << "\n";
11    up1.reset();
12    std::cout << *op1 << "\n";
13 }
```

observer.cpp

Removing New/Delete

We can use another function to remove the need for the new keyword. It has other benefits that we will explore later.

```
1 #include <iostream>
2 #include <memory>
3
4 auto main() -> int {
5     // 1 - Worst - you can accidentally own the resource multiple
6     // times, or easily forget to own it.
7     // auto* silly_string = new std::string{"Hi"};
8     // auto up1 = std::unique_ptr<std::string>(silly_string);
9     // auto up11 = std::unique_ptr<std::string>(silly_string);
10
11     // 2 - Not good - requires actual thinking about whether there's a leak.
12     auto up2 = std::unique_ptr<std::string>(new std::string("Hello"));
13
14     // 3 - Good - no thinking required.
15     auto up3 = std::make_unique<std::string>("Hello");
16
17     std::cout << *up2 << "\n";
18     std::cout << *up3 << "\n";
19     // std::cout << *(up3.get()) << "\n";
20     // std::cout << up3->size();
21 }
```

smart-no-new.cpp



Shared Pointer

- **`std::shared_ptr<T>`**
- Several shared pointers share ownership of the object
 - A reference counted pointer
 - When a shared pointer is destructed, **if it is the only shared pointer left pointing at the object, then the object is destroyed**
 - May also have many observers
 - Just because the pointer has shared ownership doesn't mean the observers should get ownership too
- **`std::weak_ptr<T>`**
 - Weak pointers are used with shared pointers when:
 - You don't want to add to the reference count
 - You want to be able to check if the underlying data is still valid before using it



Shared Pointer

```
1 #include <iostream>
2 #include <memory>
3
4 auto main() -> int {
5     auto x = std::make_shared<int>(5);
6     auto y = std::shared_ptr<int>(x);
7
8     std::cout << "use count: " << x.use_count() << "\n";
9     std::cout << "value: " << *x << "\n";
10    x.reset(); // Memory still exists, due to y.
11    std::cout << "use count: " << y.use_count() << "\n";
12    std::cout << "value: " << *y << "\n";
13    y.reset(); // Deletes the memory, since
14    // no one else owns the memory
15    std::cout << "use count: " << x.use_count() << "\n";
16    std::cout << "value: " << *y << "\n";
17 }
```

shared.cpp

Weak Pointer

```
1 #include <iostream>
2 #include <memory>
3
4 auto main() -> int {
5     auto x = std::make_shared<int>(1);
6
7     auto wp = std::weak_ptr<int>(x); // x owns the memory
8
9     auto y = wp.lock();
10    if (y != nullptr) { // x and y own the memory
11        // Do something with y
12        std::cout << "Attempt 1: " << *y << '\n';
13    }
14 }
```

weak.cpp



When To Use Which

- **Unique pointer vs Shared pointer**
 - You almost always want a unique pointer over a shared pointer
 - Use a shared pointer if either:
 - An object has multiple owners, and **you don't know which one will stay around the longest**
 - You need temporary ownership (unlikely)



Examples Of Smart Pointer Usage

- Linked list
- Doubly linked list
- Tree
- Graph



Leak Freedom

“Leak freedom in C++” poster

Strategy	Natural examples	Cost	Rough frequency
1. Prefer scoped lifetime by default (locals, members)	Local and member objects – directly owned	Zero: Tied directly to another lifetime	O(80%) of objects
2. Else prefer <code>make_unique</code> & <code>unique_ptr</code> or a container , if the object must have its own lifetime (i.e., heap) and ownership can be unique w/o owning cycles	Implementations of trees, lists	Same as new/delete & malloc/free Automates simple heap use in a library	
3. Else prefer <code>make_shared</code> & <code>shared_ptr</code> , if the object must have its own lifetime (i.e., heap) and shared ownership w/o owning cycles	Node-based DAGs, incl. trees that share out references	Same as manual reference counting (RC) Automates shared object use in a library	O(20%) of objects

Don't use owning raw *'s == don't use explicit *delete*

Don't create ownership cycles across modules by owning “upward” (violates layering)

Use *weak_ptr* to break cycles



Stack Unwinding

- Stack unwinding is the process of exiting the stack frames until we find an exception handler for the function
- This calls any destructors on the way out
 - Any resources not managed by destructors won't get freed up
 - If an exception is thrown during stack unwinding, `std::terminate` is called



Stack Unwinding

```
1 void g() {
2   throw std::runtime_error{""};
3 }
4
5 int main() {
6   auto ptr = new int{5};
7   g();
8   // Never executed.
9   delete ptr;
10 }
```

```
1 void g() {
2   throw std::runtime_error{""};
3 }
4
5 int main() {
6   auto ptr = new int{5};
7   g();
8   auto uni = std::unique_ptr<int>(ptr);
9 }
```

```
1 void g() {
2   throw std::runtime_error{""};
3 }
4
5 int main() {
6   auto ptr = std::make_unique<int>(5);
7   g();
8 }
```



Exceptions And Destructors

- During stack unwinding, `std::terminate()` will be called if an exception leaves a destructor
- The resources may not be released properly if an exception leaves a destructor
- All exceptions that occur inside a destructor should be handled inside the destructor
- Destructors usually don't throw, and need to explicitly opt in to throwing
 - STL types don't do that



Partial Construction

- What happens if an exception is thrown halfway through a constructor?
 - The C++ standard: "An object that is partially constructed or partially destroyed will have destructors executed for all of its fully constructed subobjects"
 - A destructor is not called for an object that was partially constructed
 - Except for an exception thrown in a constructor that delegates (why?)

```
1 #include <exception>
2
3 class my_int {
4 public:
5     my_int(int const i) : i_{i} {
6         (void)i_;
7         if (i == 2) {
8             throw std::exception();
9         }
10    }
11 private:
12    int i_;
13 };
14
15 class unsafe_class {
16 public:
17     unsafe_class(int a, int b)
18         : a_{new my_int{a}}
19         , b_{new my_int{b}}
20         {}
21
22     ~unsafe_class() {
23         delete a_;
24         delete b_;
25     }
26 private:
27     my_int* a_;
28     my_int* b_;
29 };
30
31 int main() {
32     auto a = unsafe_class(1, 2);
33 }
```

partial-construction-bad.cpp



Partial Construction: Solution

- Option 1: Try / catch in the constructor
 - Very messy, but works (if you get it right...)
 - Doesn't work with initialiser lists (needs to be in the body)
- Option 2:
 - An object managing a resource should initialise the resource last
 - The resource is only initialised when the whole object is
 - Consequence: An object can only manage one resource
 - If you want to manage multiple resources, instead manage several wrappers , which each manage one resource

```
1 #include <exception>
2 #include <memory>
3
4 class my_int {
5 public:
6     my_int(int const i)
7     : i_{i} {
8         (void)i_;
9         if (i == 2) {
10             throw std::exception();
11         }
12     }
13 private:
14     int i_;
15 };
16
17 class safe_class {
18 public:
19     safe_class(int a, int b)
20     : a_(std::make_unique<my_int>(a))
21     , b_(std::make_unique<my_int>(b))
22     {}
23 private:
24     std::unique_ptr<my_int> a_;
25     std::unique_ptr<my_int> b_;
26 };
27
28 int main() {
29     auto a = safe_class(1, 2);
30 }
```

partial-construction-good.cpp

👂 Feedback



Or go to the [form here](#).

