Thoughts on Assignment 2 Error Handling Using Exceptions

CS 311 Data Structures and Algorithms Lecture Slides Friday, September 13, 2024

Glenn G. Chappell Department of Computer Science University of Alaska Fairbanks ggchappell@alaska.edu © 2005-2024 Glenn G. Chappell Some material contributed by Chris Hartman

### Topics: Advanced C++

- ✓ Expressions
- ✓ Parameter passing I
- Operator overloading
- Example class
- ✓ Parameter passing II
- Invisible functions I
- Managing resources in a class
- Containers & iterators
- ✓ Invisible functions II
  - Error handling
  - Using exceptions

## Topics: S.E. Concepts

- ✓ Abstraction
- ✓ Assertions
- Testing
- ✓ Invariants

## Review

An **invariant** is a condition that is always true at a particular point in a computation. Example: the condition used in an assertion.

#### Three Special Kinds

- Precondition. Invariant at the beginning of a function. The responsibility for ensuring the preconditions hold lies with the caller.
  - What must be true for the function to execute properly.

Pre- and postconditions are the basis for **operation contracts**.

- Postcondition. Invariant at the end of a function. Tells what services the function has performed.
  - Describe the function's effect using statements about values.
- Class invariant. Invariant that holds whenever an object of the class exists, and execution is not inside a member function.
  - Statements about data members that indicate what it means for an object to be valid or usable.
  - These are preconditions for all member functions except ctors, and postconditions for all member functions except the dctor.

#### Review Software Engineering Concepts: Invariants — TRY IT (Exercise)

## Exercise

 Write pre- and postconditions for the one-parameter constructor for class Abc.

See next slide for answers.

// class Abc // Invariants: // 0 <= n && n < 100 class Abc { public: Abc(int nn) : n(nn) { } [other stuff here] private: int n; }; // End class Abc

#### Review Software Engineering Concepts: Invariants — TRY IT (Answers)

## Exercise

 Write pre- and postconditions for the one-parameter constructor for class Abc.

#### Answers

// Pre:

0 <= \_n && \_n < 100
is also a postcondition.
But that is already in the
class invariants, and we
do not need to repeat it.</pre>

// class Abc // Invariants: // 0 <= n && n < 100 class Abc { public: Abc(int nn) : n(nn) { } [other stuff here] private: int n; }; // End class Abc

# Thoughts on Assignment 2

In Assignment 2 you write a "moderately smart" array (MSArray). This will require applying some recently covered ideas.

- Managing Resources in a Class
  - Are you doing dynamic allocation correctly? When you allocate something, is it always freed?
  - MSArray uses RAII.
  - What type to use for the size of an MSArray? For array indices?
- Containers & Iterators
  - MSArray is a generic container. Its member functions begin and end return iterators.
- Software Engineering Concepts: Invariants
  - You will need to document preconditions for each function that has them and class invariants for each class.
- Invisible Functions II
  - MSArray directly manages a resource. You will need to define all of the Big Five. (Do not apply the Rule of Zero in Assignment 2!)

MSArray is a class template, like std::vector.

MSArray<int> ia;



Templates go entirely in the header. Do not write a separate source file for MSArray.

Each associated global function is actually a **function template**. Define one of these (in the header file!) like this:



Document **class invariants** for all classes and **preconditions** for all functions that have them.

```
// Invariants: ...
template <typename ValType>
class MSArray {
    ...
};
                  If there are any
                  preconditions.
// Pre: ...
template <typename ValType>
bool operator==( ... )
```

Member functions begin and end return iterators.

- *These can be pointers.* Do not write a separate iterator class.
- Function begin returns an iterator to the first array item. You already have a pointer to the first array item (*think ...*); use it.
- Function end returns an iterator to just past the last array item. Add a number to the return value of begin (what number?).



A const MSArray has non-modifiable data. If a function can give access to data in modifiable form, then write two versions.

```
... operator[]( ... )
{ ... }
const ... operator[]( ... ) const
{ ... }
```

```
... begin()
{ ... }
const ... begin() const
{ ... }
```

In each pair, the two functions should be identical, except for (1) the const at the end of the first line, and (2) what is returned—or how it is returned.

In this particular case, we will allow repetition of code.

```
... end()
```

•••

You will need to write the Big Five in MSArray.

Items in C++ built-in arrays are always default-constructed. We cannot set their values to anything else in a member initializer. Therefore, the copy ctor will need a loop\* in the function body.



For the rest, see Invisible Functions II, and do what it says!

CS 311 Fall 2024

# **Error Handling**

An **error condition** (often *error*) is a condition occurring during runtime that cannot be handled by the normal flow of execution.

An error condition is not the same as a bug in the code.

- We are not referring to compilation errors.
- Some error conditions are caused by bugs, but our discussion of error handling will focus on properly written code.

An error condition does not mean the user did something wrong.

Some error conditions are caused by user mistakes.

Example

- A function copyFile opens a file, reads its contents, and writes them to another file.
- copyFile is called to read a file that is accessed via a network.
- Halfway through reading the file, the network goes down.
- It is now impossible to read the file. The normal flow of execution cannot handle this situation. We have an error condition.

How do we deal with possible error conditions?

Sometimes we can **prevent error conditions**:

- Write a precondition that requires the caller to keep a certain problem from happening.
- Example. Insisting on a non-zero parameter, to prevent a division-by-zero error condition.

## Sometimes we can contain error conditions,

by handling them ourselves:

- If something is not right, then deal with it.
- Example. A fast algorithm needs more memory than we have; we use a slow method instead.

But sometimes neither of these two is feasible.

Then we must **signal the client code**.

- Signal the client code when the function is unable to fulfill its postconditions.
- Example. The earlier file-reading troubles.

Handle a possible error condition **before** the function.

Handle a possible error condition **in** the function.

Handle a possible error condition **after** the function.

- When client code might need to be informed of an error condition, we may have these three goals:
- Error conditions must not wreck our program. It must continue running, and later end properly. Objects must be usable. Resources must not leak.
- Even better, it would be good if each operation we attempt either completes successfully, or, if there is an error condition, has no effect.
- But it would be really great if we never need to inform client code of errors at all.

# Later in the class, we will formalize these as **safety guarantees**.

The first goal is the fundamental standard that all code must meet. We call it the **Basic Guarantee**.

The second is preferred, although sometimes not feasible. We call it the **Strong Guarantee**.

The third is often wishful thinking. Sometimes we simply *must* inform client code of an error condition. But in special cases—often involving *finishing* something—we do require this standard. We call it the **No-Throw Guarantee** (some call it the **No-Fail Guarantee**). When we cannot prevent or contain an error condition, then we must signal the client code. How can we do this?

Method 1. Return an error code.

```
int c = getc(myFile);
if (c == EOF)
    printf("End of file\n");
```

The old C-language I/O library uses this method.

Method 2. Set a flag to be checked by a separate error-checking function.

char c;

myFileStream >> c;

C++ file streams default to using this method.

if (myFileStream.eof())

cout << "End of file" << endl;</pre>

2024-09-13

CS 311 Fall 2024

Return codes and separate error-checking functions are acceptable methods for flagging error conditions, but they have downsides.

- They can be difficult to use in places where a value cannot be returned, or an error condition cannot be checked for: constructors, in the middle of an expression, etc.
- They can lead to complicated code.

Because of these issues, another method was developed.

Method 3. Throw an exception.

Exceptions are available in many programming languages: C++, Java, Python, JavaScript, etc. They are associated with OOP.

In our next topic, we will look at exceptions in C++.

### Error Handling Summary

An **error condition** (often *error*) is a condition occurring during runtime that cannot be handled by the normal flow of execution.

- Not necessarily a bug or a user mistake.
- Example: Could not read file.



# **Using Exceptions**

**Exception**: an object that is **thrown** to signal an error condition.

new throws std::bad\_alloc or a derived class, if allocation fails.
To handle an exception, catch it using try ... catch.

```
#include <new> // for std::bad alloc
                                                    catch gets an exception
                                                    that is thrown inside the
Foo * p;
                                                    corresponding try-block,
                                                    if it has the proper type.
bool success = true; e is the exception.
try {
     p = new Foo;
                                                     Standard exception types
                                   Catch exceptions
                                                     have a member function
                                   by reference.
}
                                                     what. It returns string.
catch (std::bad alloc & e)
     success = false;
     cerr << "Allocation failed: " << e.what() << endl;
```

Under what circumstances is a thrown exception caught? If it is caught, then *where* in the code is it caught?

#### How It Works

- When an exception is thrown inside a try-block, control immediately passes to the catch-block that is associated with the smallest enclosing try-block that catches the proper type. Derived classes are handled as usual.
- In all other circumstances, a catch-block is not executed.
- An uncaught exception terminates the program.

That's it! Exception handling is not complicated—even if some of the examples we cover make it seem complicated.

A catch only gets an exception that is:

- Thrown inside the corresponding try-block.
- Of an appropriate type.

Once an exception is thrown, the try-block is exited. If no exception is thrown, the catch-block is not executed.



catch gets exceptions of the proper type that are thrown inside the corresponding try.

This includes an exception thrown in a called function, if it is not caught inside that function—that is, if it **escapes** the function.

```
void myFunc()
{
                                          Function main would be able to catch an
     globalP1 = new Foo; 
                                          exception thrown by this statement ...
     globalFlag = true;
     try {
          qlobalP2 = new Foo; ←
                                                 ... but not a std::bad alloc
                                                 thrown by this statement.
     catch (std::bad alloc & e) {
          globalFlag = false;
```

Exceptions can propagate out of nested function calls. Catching by reference will catch exceptions of derived types.

void xx(); // May throw std::bad alloc void yy() When function *zz* is called, if { XX(); } function xx throws std::bad alloc, then the exception will be caught here. void zz() { try { yy(); } catch (std::exception & e) {

Because we catch by reference, derived classes of std::exception will be caught.

All standard exception classes, including std::bad\_alloc, are derived from std::exception.

### The following do *not* throw:

- Built-in operations, other than new, on built-in types.
  - Including operator[].
- Deallocation done by the built-in version of delete.
  - Note. delete calls destructors, which conceivably might throw—but should not, as we will see.
- C++ Standard I/O Libraries (default behavior).

### The following can throw:

- new may throw std::bad\_alloc or a derived class (default behavior).
- A function that (1) calls a function that throws, and (2) does not catch the exception, will throw.
- Functions written by others may throw. See their documentation.
- Lastly, throw always throws. See the next subtopic.

We can throw our own exceptions, using throw.

```
We do not do
class CC {
                                            this very much!
public:
    int & operator[](std::size t ix)
         // May throw std::out of range
     {
         if (ix >= arrsize)
              throw std::out of range("CC::op[]: bad ix");
         return arr[ix];
     }
                                  The syntax of throw is just like
                                  the syntax of return.
private:
    int *
                  arr;
    std::size t arrsize;
```

When throwing your own exception—which you will not do very much!—use or derive from one of the standard exception types. Standard exception types are set up to allow for derived classes; in particular, they all have virtual destructors.

Standard exception types have a string member, for a humanreadable message. This is a ctor parameter. Access it via the what() member function. Use catch(...) to catch all exceptions. Inside a catch-block, "throw;" will re-throw the same exception. These two are used together, to ensure that clean-up gets done.



Now we know two ways to ensure that clean-up is done before we leave: (1) RAII, (2) catch all & re-throw.

- **Fact 1.** An automatic object's dctor is called when it goes out of scope, even if this is due to an exception.
- **Fact 2.** If an exception is thrown, and one of the destructors called before it is caught also throws, then the program terminates.

Put these two facts together, and we conclude:

**Destructors should not throw.** 

The above is a technical argument based on the specification of C++. From a more philosophical point of view, *finishing-up* operations—like destructors—generally should not throw.

2024-09-13

Dctors are only called for **fully constructed** objects. If a ctor throws, then the dctor for that object will not be called.

> It is okay for constructors to throw.

Because dctors should not throw, they are generally marked noexcept implicitly, unless otherwise specified.

If a noexcept function throws, then the program terminates.

Recall: noexcept is a promise that a function will not throw.

We can make a destructor that is not noexcept using "noexcept(false)". However, this is EVIL. 😕



### Using Exceptions CODE

## TO DO

• Write a function allocate2 that:

Done. See allocate2.cpp.

- Attempts to allocate two dynamic objects.
- Returns pointers to these objects, using reference parameters.
- If either allocation fails, throws std::bad\_alloc.
- Has no memory leaks.
- Look at example code showing how RAII can simplify these situations.

See allocate2\_raii.cpp.

When to Do Things

- Catch when you can handle an error condition that may be signaled by some function you call.
- Throw when your function is unable to fulfill its postconditions and must signal an error condition.
- Catch all & re-throw when you call a function that may throw, you cannot handle the error, but you do need to do some clean-up before your function exits.

**Typically we do not write more than one of these.** If someone writes a throw, the catch is typically written by someone else.



Some people do not like exceptions. Some of these people are very vocal about their dislike. But I think that *some* of them dislike exceptions for the wrong reasons.

A bad reason to dislike exceptions is that they require lots of work.

- Dealing with error conditions is work. Writing software that works is work. Exceptions are one tool we can use to achieve this goal.
- Handling exceptions properly is hard work because writing correct software is hard work.

What *might* be a good reason to dislike exceptions is that they add hidden execution paths. But remember that other error-handling methods have their own downsides—which is why exceptions were invented.