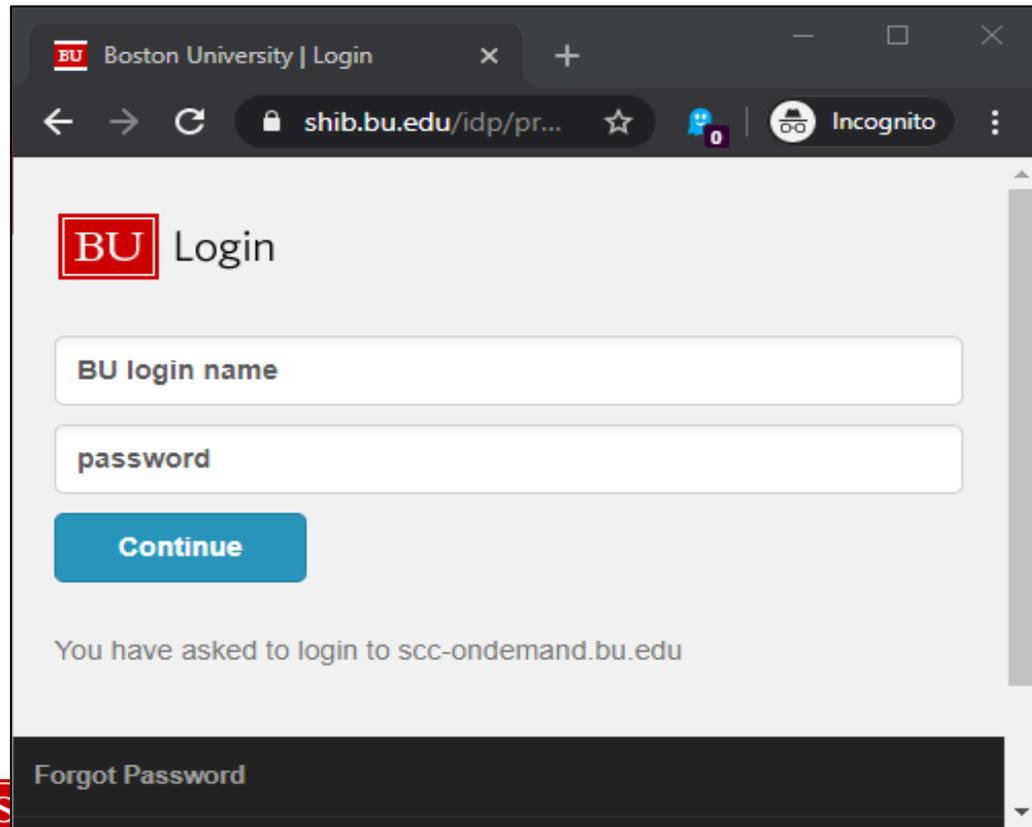


Introduction to C++: Part 4

Existing SCC Account

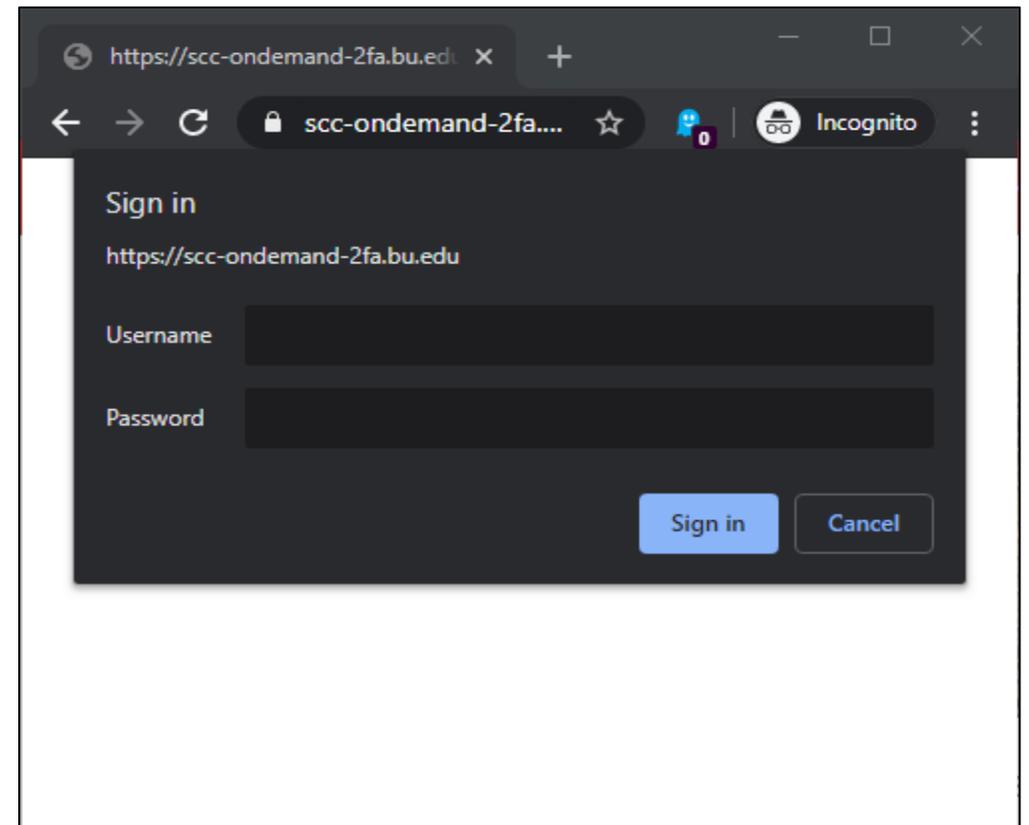
1. Open a web browser
2. Navigate to <http://scc-ondemand.bu.edu>
3. Log in with your BU Kerberos Credentials



The screenshot shows a web browser window with the address bar displaying "shib.bu.edu/idp/pr...". The page content includes the Boston University logo and the text "Login". Below this, there are two input fields: "BU login name" and "password". A blue "Continue" button is positioned below the password field. At the bottom of the page, there is a "Forgot Password" link. A message at the bottom of the page reads "You have asked to login to scc-ondemand.bu.edu".

Temporary Tutorial Account

1. Open a web browser
2. Navigate to <http://scc-ondemand-tutorial.bu.edu>
3. Log in with Tutorial Account



The screenshot shows a login dialog box in a web browser window. The address bar displays "https://scc-ondemand-2fa.bu.edu". The dialog box has a dark background and contains the text "Sign in" and "https://scc-ondemand-2fa.bu.edu". Below this, there are two input fields: "Username" and "Password". At the bottom right of the dialog, there are two buttons: "Sign in" and "Cancel".

Click on Interactive Apps/Desktop



- Desktops
 - Desktop
 - MATLAB
 - Mathematica
 - QGIS
 - SAS
 - STATA
 - Spyder
 - VirtualGL Desktop
- Servers
 - Jupyter Notebook
 - RStudio Server
 - Shiny App Server
 - TensorBoard Server



Access the SCC using only your web browser!

[SCC OnDemand Documentation](#)

- Interactive Apps
- Desktops
 - Desktop**
 - MATLAB
 - Mathematica
 - QGIS
 - SAS
 - STATA
 - Spyder
 - VirtualGL Desktop
- Servers
 - Jupyter Notebook
 - RStudio Server
 - Shiny App Server
 - TensorBoard Server
 - Webserver

Desktop

This app will launch an interactive desktop on a compute node.

List of modules to load (space separated)

eclipse/2019-06 gcc/8.3.0 Select Modules

eclipse/2019-06
gcc/8.3.0

Working Directory

Select Directory

The directory to start in. (Defaults to home directory.)

Initial command to run

xfce4-terminal

Number of hours

3

3

Number of cores

1

Number of gpus

0

Project

scv

Extra qsub options

I would like to receive an email when the session starts

Launch

click

* The Desktop session data for this session can be accessed under the [data root directory](#).



Desktop (6924) 1 core | Running

Host: [>_scc-w12](#) Delete

Created at: 2020-02-04 14:53:50 EST

Time Remaining: 2 hours and 59 minutes

Session ID: 41466d74-9ac7-4f79-b596-26cffdf6cf9b

Compression Image Quality

0 (low) to 9 (high) 0 (low) to 9 (high)

[Connect to Desktop](#) View Only (Share-able Link)

When your desktop is ready click *Connect to Desktop*

- Enter this command to create a directory in your home folder and to copy in tutorial files:

```
/net/scc2/scratch/intro_to_cpp4.sh
```

or

Download Part4.zip:

```
http://rcs.bu.edu/examples/cpp/tutorial/
```

C++ Libraries

- There are a ***LOT*** of libraries available for C++ code.
 - [Sourceforge](#) alone has > 9400
- Before jumping into writing your code, consider what you need and see if there are libraries available.
- Many libraries contain code developed by professionals or experts in a particular field.
- Consider what you are trying to accomplish in your research:
 - A) accomplishments in your field or
 - B) C++ programming?

C++ Compilers on the SCC

Module name	Vendor	Compiler	Versions
gnu	GNU	g++	4.8.5 - 11.2.0
intel	Intel	icpc	2016 - 2021.1
pgi	Portland Group / Nvidia	pgc++	16.5 - 19.4
llvm	LLVM	clang++	3.9 .1- 12.0.1

- There are 4 families of compilers on the SCC for C++.
 - To see versions use the *module avail* command, e.g. `module avail gnu`
- They have their strengths and weaknesses.
- For info on how to choose compiler optimizations for the SCC see the RCS website:
<http://www.bu.edu/tech/support/research/software-and-programming/programming/compilers/compiler-optimizations/>

C++ Standard by Compiler

- [Support for C++ standards in g++](#)
 - Intel icpc: On Linux, g++ header files are used by the Intel icpc compiler, so icpc will support the standards used by the available g++.
- [Support in Microsoft Visual C++ compiler](#)
- [Support in clang++](#)
 - (as used on Mac OSX)

Multithreading

- OpenMP
 - Open MP is a standard approach to writing multithreaded code to exploit multiple CPU cores with your program.
 - Fully supported in C++
- Intel Thread Building Blocks
 - C++ specific library
 - Available on the SCC from Intel and is also open source. (in the intel modules)
 - Much more flexible and much more C++-ish than OpenMP
 - Offers high performance memory allocators for multithreaded code
 - Includes concurrent data types (vectors, etc.) that can automatically be shared amongst threads with no added effort for the programmer to control access to them.
- Data Parallel C++
 - Dialect of C++ with extensive multi-threading built in.



Math and Linear Algebra

- Eigen
 - http://eigen.tuxfamily.org/index.php?title=Main_Page
 - “Eigen is a C++ template library for linear algebra: matrices, vectors, numerical solvers, and related algorithms.”
- Armadillo
 - <http://arma.sourceforge.net/>
 - “Armadillo is a high quality linear algebra library (matrix maths) for the C++ language, aiming towards a good balance between speed and ease of use. Provides high-level syntax (API) deliberately similar to Matlab.”
- OpenCV
 - A computer vision and image processing library, with excellent high-performance support for linear algebra, many algorithms, and GPU acceleration.
- Ceres
 - non-linear optimization
- LAPACK++
 - C++ wrapper for the BLAS and LAPACK libraries
- dlib
 - Machine learning and data analysis

Other useful libraries

- Parsers

- CLI11 - <https://github.com/CLIUtils/CLI11>
 - Command line arguments. Header-only library, C++11 standard required
- json - <https://github.com/nlohmann/json>
 - JSON format reading/writing. Header-only library.

- Physical units (enforced at compile time!):

- mp-units - <https://github.com/mpusz/units>
 - In consideration to be included in the C++23/26 standard.
 - Needs C++20 to compile and use
- units - <https://github.com/nholthaus/units>
 - Header-only library, requires C++14

- Random Numbers

- [C++11](#) standard RNGs
- [PCG](#) library
 - Faster number generation, works with C++11 RNG containers



The doomed [Mars Observer](#) spacecraft.

Using subclasses

- A function that takes a superclass argument can *also* be called with a subclass as the argument.
- The reverse is **not** true – a function expecting a subclass argument cannot accept its superclass.
- Copy the code to the right and add it to your main.cpp file.

```
void PrintArea (Rectangle &rT) {  
    cout << rT.Area () << endl ;  
}  
  
int main () {  
    Rectangle rT (1.0, 2.0) ;  
    Square sQ (3.0) ;  
    PrintArea (rT) ;  
    PrintArea (sQ) ;  
}
```

The PrintArea function can accept the Square object sQ because Square is a subclass of Rectangle.



Overriding Methods

- Sometimes a subclass needs to have the same interface to a method as a superclass but with different functionality.
- This is achieved by *overriding* a method.
- Overriding a method is simple: just re-implement the method with the same name and arguments in the subclass.

```
class Super {
public:
    void PrintNum() {
        cout << 1 << endl ;
    }
};

class Sub : public Super {
public:
    // Override
    void PrintNum() {
        cout << 2 << endl ;
    }
};

Super sP ;
sP.PrintNum() ; // Prints 1
Sub sB ;
sB.PrintNum() ; // Prints 2
```



Overriding Methods

- Seems simple, right?

```
class Super {
public:
    void PrintNum() {
        cout << 1 << endl ;
    }
};

class Sub : public Super {
public:
    // Override
    void PrintNum() {
        cout << 2 << endl ;
    }
};

Super sP ;
sP.PrintNum() ; // Prints 1
Sub sB ;
sB.PrintNum() ; // Prints 2
```

How about in a function call...

- Using a single function to operate on different types is *polymorphism*.
- Given the class definitions, what is happening in this function call?

“C++ is an insult to the human brain”
– Niklaus Wirth (designer of Pascal)

```
class Super {
public:
    void PrintNum() {
        cout << 1 << endl ;
    }
};

class Sub : public Super {
public:
    // Override
    void PrintNum() {
        cout << 2 << endl ;
    }
};
```

```
void FuncRef (Super &sP) {
    sP.PrintNum() ;
}

Super sP ;
Func (sP) ; // Prints 1
Sub sB ;
Func (sB) ; // Hey!! Prints 1!!
```

Type casting

```
void FuncRef (Super &sP) {  
    sP.PrintNum() ;  
}
```



- The Func function passes the argument as a *reference* (Super &sP).
 - What's happening here is *dynamic type casting*, the process of converting from one type to another at runtime.
 - Same mechanism as the *dynamic_cast<type>()* function
- The incoming object is treated as though it were a superclass object in the function.
- When methods are overridden and called there are two points where the proper version of the method can be identified: either at compile time or at runtime.

Virtual methods

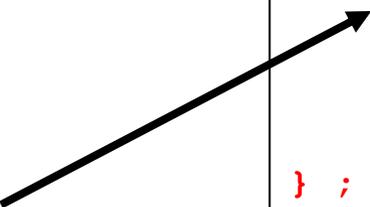
- When a method is labeled as virtual and overridden the compiler will generate code that will check the type of an object at **runtime** when the method is called.
- The type check will then result in the expected version of the method being called.
- When overriding a virtual method in a subclass, it's a good idea to label the method as virtual in the subclass as well.
 - ...just in case this gets subclassed again!

```
class SuperVirtual
{
public:
    virtual void PrintNum()
    {
        cout << 1 << endl ;
    }
};

class SubVirtual : public SuperVirtual
{
public:
    // Override
    virtual void PrintNum()
    {
        cout << 2 << endl ;
    }
};

void Func(SuperVirtual &sP)
{
    sP.PrintNum() ;
}

SuperVirtual sP ;
Func(sP) ; // Prints 1
SubVirtual sB ;
Func(sB) ; // Prints 2!!
```

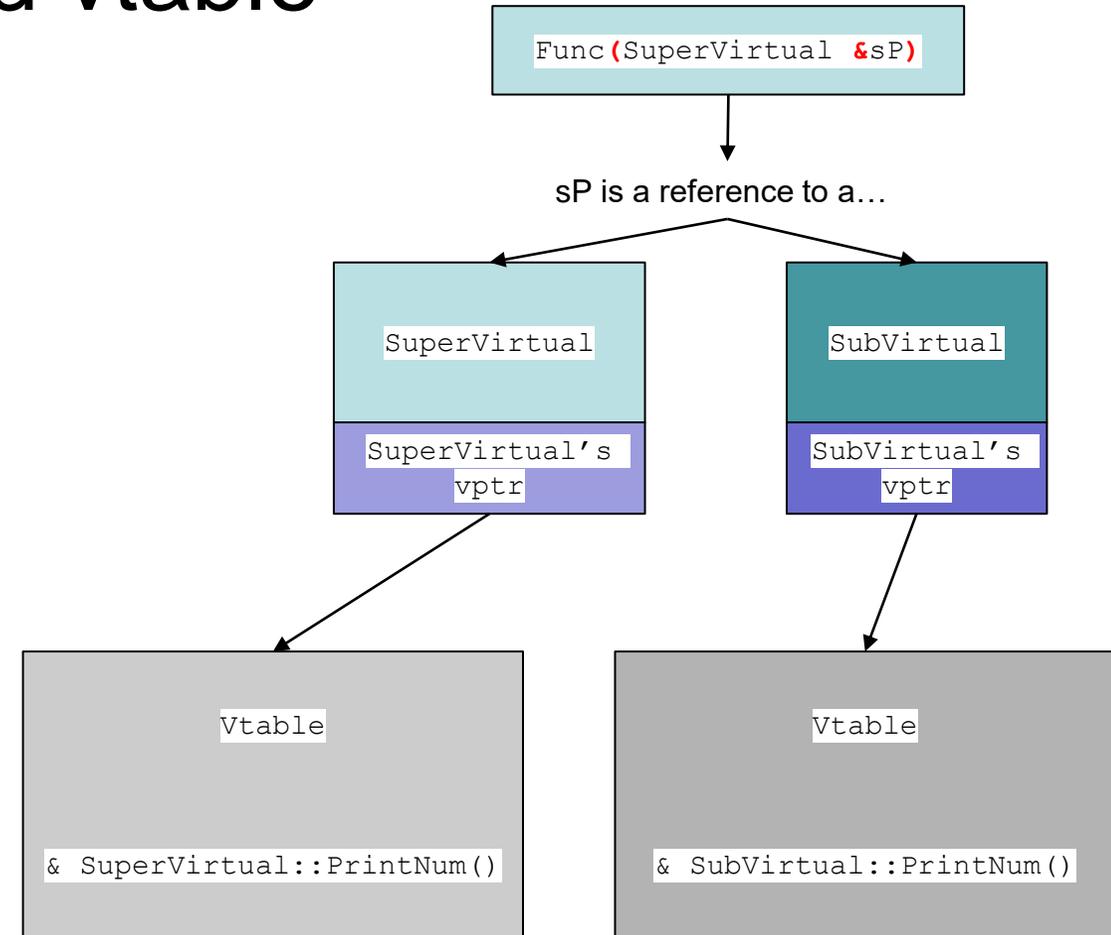


Early (static) vs. Late (dynamic) binding

- Leaving out the virtual keyword on a method that is overridden results in the compiler deciding *at compile time* which version (subclass or superclass) of the method to call.
- This is called early or static *binding*.
- At compile time, a function that takes a superclass argument will only call the **non-virtual** superclass method under early binding.
- Making a method virtual adds code behind the scenes (that you, the programmer, never interact with directly)
 - Lookups in a hidden table, called the *vtable*, are done to figure out what version of the virtual method should be run.
- This is called late or dynamic binding.
- There is a small performance penalty for late binding due to the vtable lookup.
- **This only applies when an object is referred to by a reference or pointer.**

Behind the scenes – vptr and vtable

- C++ classes have a hidden pointer (vptr) generated that points to a table of virtual methods associated with a class (vtable).
- When a virtual class method (base class or its subclasses) is called by reference (or pointer) *when the program is running* the following happens:
 - The object's **class** vptr is followed to its **class** vtable
 - The virtual method is looked up in the vtable and is then called.
 - One vptr and one vtable per class so minimal memory overhead
 - If a method override is **non-virtual** it won't be in the vtable and it is selected at **compile time**.



Let's run this through the debugger

- Open the project `Virtual_Method_Calls`.
- Everything here is implemented in one big `main.cpp`
- Place a breakpoint at the first line in `main()` and in the two implementations of `Func()`



When to make methods virtual

- If a method will be (or might be) overridden in a subclass, make it virtual
 - There is a *minuscule* performance penalty. Will that even matter to you?
 - i.e. Have you profiled and tested your code to show that virtual method calls are a performance issue?
 - When is this true?
 - Almost always! Who knows how your code will be used in the future?
- Constructors are **never** virtual in C++.
- Destructors in a base class should always be virtual.
 - Also – if any method in a class is virtual, make the destructor virtual
 - These are important when dealing with objects via reference and it avoids some subtleties when manually allocating memory.

Why all this complexity?

```
void FuncEarly(SuperVirtual &sP)
{
    sP.PrintNum();
}
```

- Called by **reference** – late binding to PrintNum()

```
void FuncLate(SuperVirtual sP)
{
    sP.PrintNum();
}
```

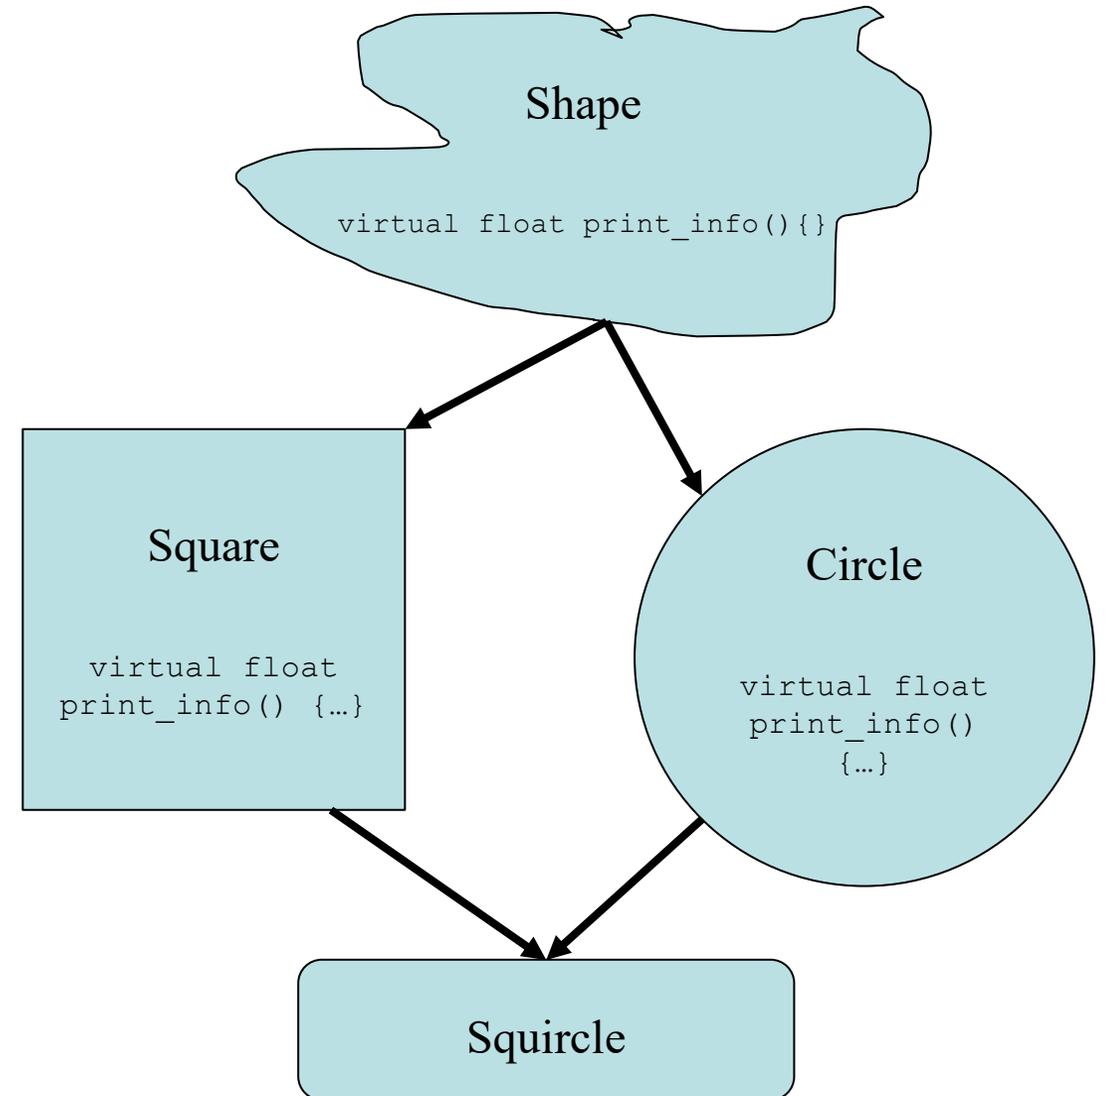
- Called by **value** – early binding to PrintNum even though it's virtual!

- Late binding allows for code libraries to be updated for new functionality. As methods are identified at runtime the executable does not need to be updated.
- This is done all the time! Your C++ code may be, for example, a plugin to an existing simulation code.
- Greater flexibility when dealing with multiple subclasses of a superclass.
- Most of the time this is the behavior you are looking for when building class hierarchies.

- Remember the Deadly Diamond of Death? Let's explain.
- Look at the class hierarchy on the right.
 - Square and Circle inherit from Shape
 - Squircle inherits from both Square and Circle
 - Syntax:

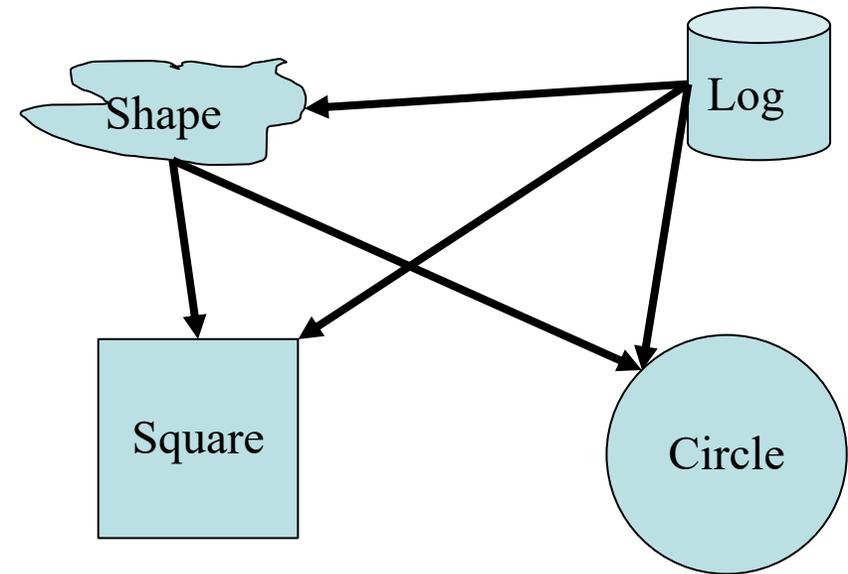
```
class Squircle : public Square, Circle
```

- The Shape class implements an empty `print_info()` method. The Square and Circle classes override it. Squircle does not.
- Under late binding, which version of `print_info()` is accessed from Squircle? Square. `info()` or Circle. `info()`?



Interfaces

- Interfaces are a way to have your classes share method names without them sharing actual code.
 - ...and hopefully methods with the same name are implemented to do the same thing, that's up to you!
- Gives much of the benefit of multiple inheritance without the complexity and pitfalls



- Example: for debugging you want each class to have a Log() method that writes some info to a file.
 - Implement with an interface.

Interfaces

- An interface class in C++ is called a pure virtual class.
- It contains virtual methods only with a special syntax. Instead of {} the function is set to 0.
 - Any subclass **must** implement pure virtual methods!
- Modified Square.h shown.
- What happens when this is compiled?

```
(...error...)
include/square.h:10:7: note:   because the following virtual
functions are pure within 'Square':
  class Square : public Rectangle, Log
    ^
include/square.h:7:18: note:   virtual void Log::LogInfo()
  virtual void LogInfo()=0 ;
```

- Once the LogInfo() is uncommented it will compile.

```
#ifndef SQUARE_H
#define SQUARE_H

#include "rectangle.h"

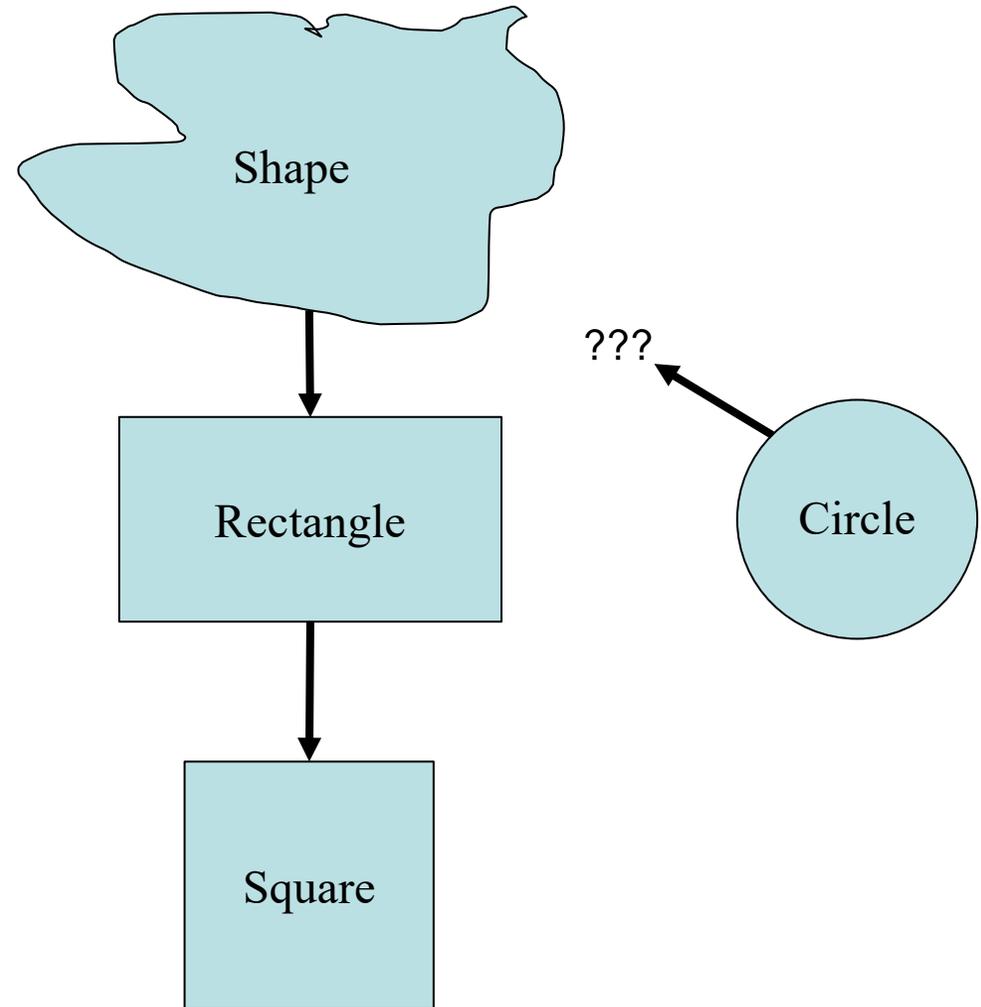
class Log {
    virtual void LogInfo()=0 ;
};

class Square : public Rectangle, Log
{
    public:
        Square(float length);
        virtual ~Square();
        // virtual void LogInfo() {}
    protected:
        private:
};

#endif // SQUARE_H
```

Putting it all together

- Now let's revisit our Shapes project.
- Open the **“Shapes with Circle”** project.
 - This has a Shape base class with a Rectangle and a Square
- Add a Circle class to the class hierarchy in a sensible fashion.



- Hint: Think first, code second.



New pure virtual Shape class

- Slight bit of trickery:
 - An empty constructor is defined in shape.h
 - No need to have an extra shape.cpp file if these functions do nothing!
- Q: How much code can be in the header file?
- A: Most of it with some exceptions.
 - .h files are not compiled into .o files so a header with a lot of code gets re-compiled every time it's referenced in a source file.
 - In other words, avoid putting source code in .h files.

```
#ifndef SHAPE_H
#define SHAPE_H

class Shape
{
    public:
        Shape () {}
        virtual ~Shape () {}

        virtual float Area ()=0 ;
    protected:

    private:
};

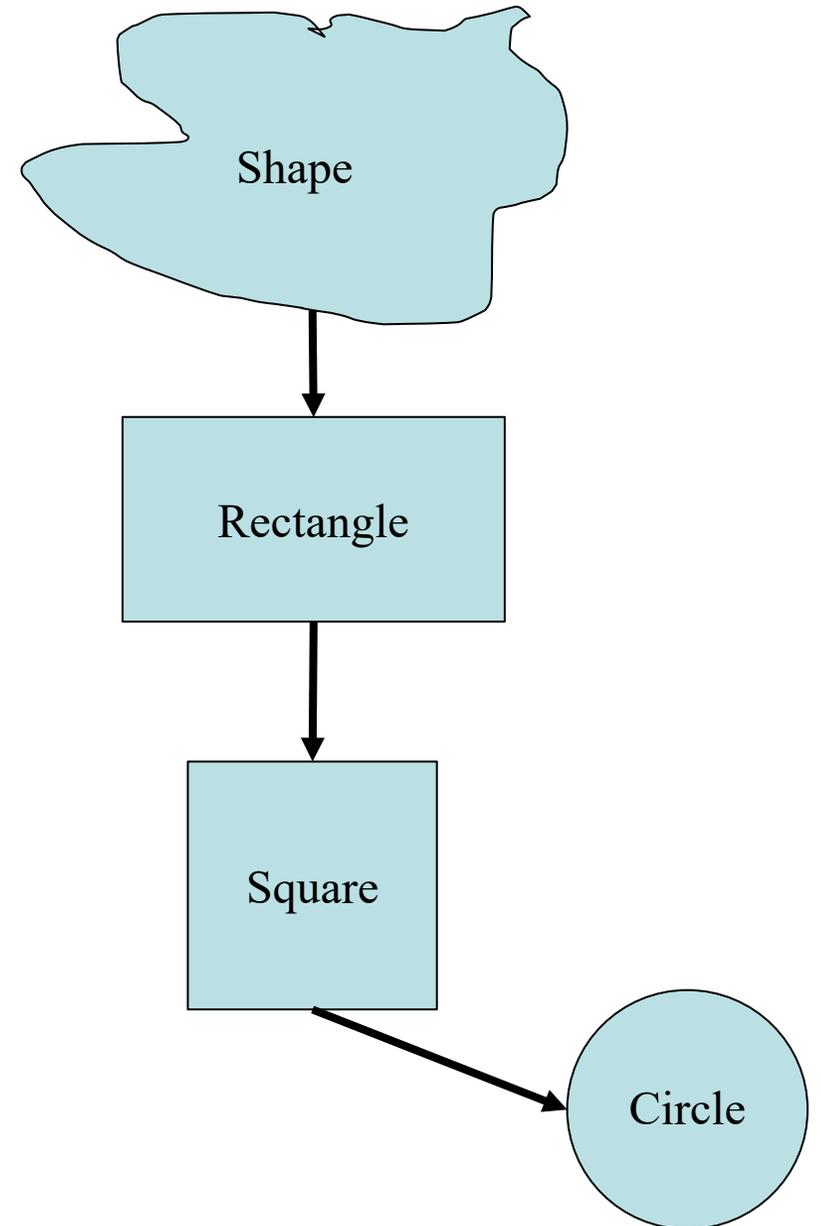
#endif // SHAPE_H
```

Give it a try

- Add inheritance from Shape to the Rectangle class
- Add a Circle class, inheriting from wherever you like.
- Implement Area() for the Circle
- If you just want to see a solution, open the project “Shapes with Circle solved”

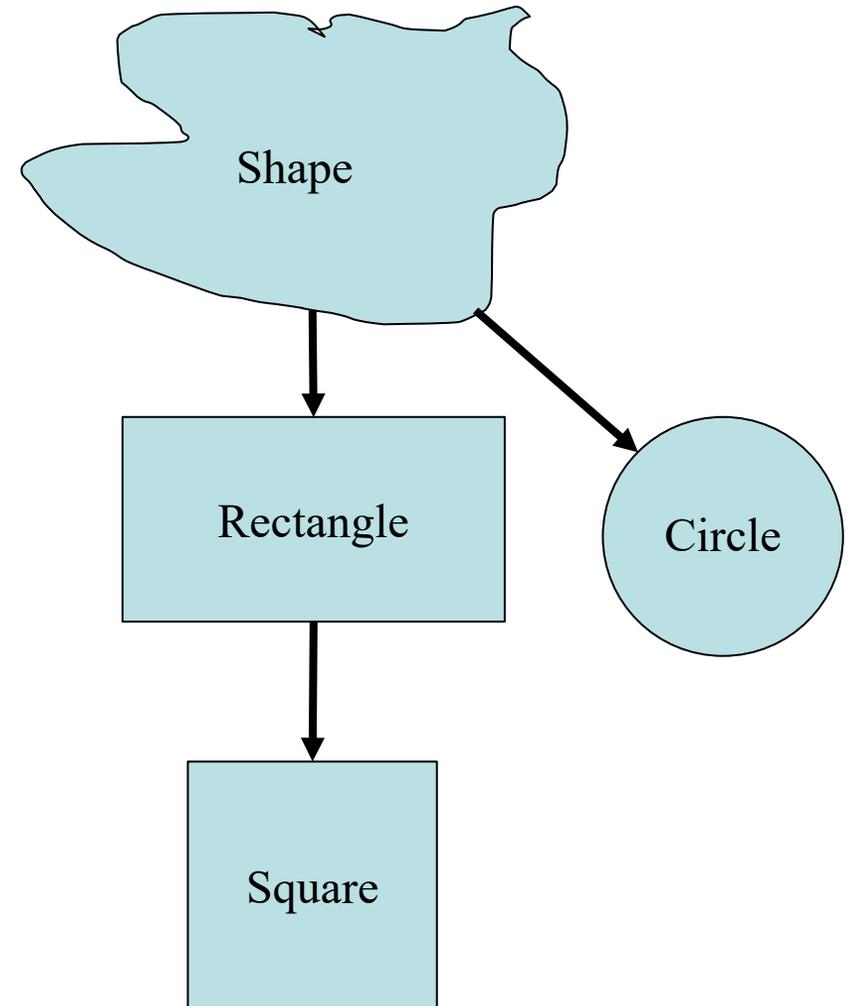
A Potential Solution

- A Circle has one dimension (radius), like a Square.
 - Would only need to override the Area() method
- But...
 - Would be storing the radius in the members m_width and m_length. This is not a very obvious to someone else who reads your code.
- Maybe:
 - Change m_width and m_length names to m_dim_1 and m_dim_2?
 - Just makes everything more muddled!



A Better Solution

- Inherit separately from the Shape base class
 - Seems logical, to most people a circle is not a specialized form of rectangle...
- Add a member `m_radius` to store the radius.
- Implement the `Area()` method
- Makes more sense!
- Easy to extend to add an Oval class, etc.



New Circle class

- Also inherits from Shape
- Adds a constant value for π
 - Constant values can be defined right in the header file.
 - If you accidentally try to change the value of PI the compiler will throw an error.

```
#ifndef CIRCLE_H
#define CIRCLE_H

#include "shape.h"

class Circle : public Shape
{
    public:
        Circle();
        Circle(float radius) ;
        virtual ~Circle();

        virtual float Area() ;

        const float PI = 3.14;
        float m_radius ;

    protected:

    private:
};

#endif // CIRCLE_H
```

- circle.cpp
- Questions?

```
#include "circle.h"

Circle::Circle()
{
    //ctor
}

Circle::~Circle()
{
    //dtor
}

// Use a member initialization list.
Circle::Circle(float radius) : m_radius{radius}
{}

float Circle::Area()
{
    // Quiz: what happens if this line is
    // uncommented and then compiled:
    //PI=3.14159 ;
    return m_radius * m_radius * PI ;
}
```

Quiz time!

- What happens behind the scenes when the function PrintArea is called?
- How about if PrintArea's argument was instead:

```
void PrintArea(Shape shape)
```

```
void PrintArea(Shape &shape) {  
    cout << "Area: " << shape.Area() << endl ;  
}  
  
int main()  
{  
    Square sQ(4) ;  
    Circle circ(3.5) ;  
    Rectangle rT(21,2) ;  
  
    // Print everything  
    PrintArea(sQ) ;  
    PrintArea(rT) ;  
    PrintArea(circ) ;  
    return 0 ;  
}
```

Quick mention...

- Aside from overriding functions it is also possible to override operators in C++.

- As seen in the C++ string. The + operator concatenates strings:

```
string str = "ABC" ;  
str = str + "DEF" ;  
// str is now "ABCDEF"
```

- It's possible to override +, -, =, <, >, brackets, parentheses, etc.

- Syntax:

```
MyClass operator*(const MyClass& mC) {...}
```

- Recommendation:

- Use with great caution. This is an easy way to write very confusing code.
 - A well-named function will almost always be easier to understand than an operator.

- An exceptions is the assignment operator: operator=

Summary

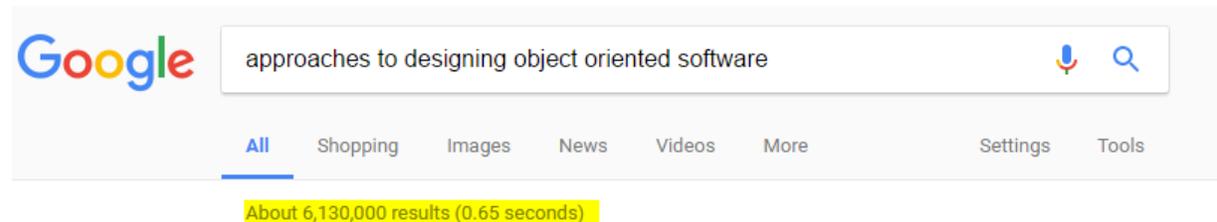
- C++ classes can be created in hierarchies via inheritance, a core concept in OOP.
- Classes that inherit from others can make use of the superclass' public and protected members and methods
 - You write less code!
- Virtual methods should be used whenever methods will be overridden in subclasses.
- Avoid multiple inheritance, use interfaces instead.
- Subclasses can override a superclass method for their own purposes and can still explicitly call the superclass method.
- Abstraction means hiding details when they don't need to be accessed by external code.
 - Reduces the chances for bugs.
- While there is a lot of complexity here – in terms of concepts, syntax, and application – keep in mind that OOP is a highly successful way of building programs!

Some OOP Guidelines

- Here are some guidelines for putting together a program using OOP to keep in mind while getting up and running with C++.
- Keep your classes simple and single purpose.
- Logically organize your classes to re-use code via inheritance.
- Use interfaces in place of multiple inheritance
- Keep your methods short
 - Many descriptive methods that do little things is easier to debug and understand.
- Follow the KISS principle:
 - “Keep it simple stupid”
 - “Keep it simple, silly”
 - “Keep it short and sweet”
 - “Make Simple Tasks Simple!” – Bjarne Stroustrup
 - “Make everything as simple as possible, but not simpler” – Albert Einstein

Putting your classes together

- Effective use of OOP demands that the programmer think/plan/design first and code second.
- There is a large body of information on this topic:



- As this is an academic institution your code may:
 - Live on in your lab long after you have graduated
 - Be worked on by multiple researchers
 - Adapted to new problems you haven't considered
 - Be shared with collaborators
- For more structured environments (ex. a team of professional programmers) there exist concepts like SOLID:
 - [https://en.wikipedia.org/wiki/SOLID_\(object-oriented_design\)](https://en.wikipedia.org/wiki/SOLID_(object-oriented_design))
 - ...and there are many others.

Keep your classes simple

- Avoid “monster” classes that implement everything including the kitchen sink.
- Our Rectangle class just holds dimensions and calculates its area.
 - It cannot print out its area, send email, draw to the screen, etc.

- **Single responsibility principle:**

- Every class has responsibility for one piece of functionality in the program.
- https://en.wikipedia.org/wiki/Single_responsibility_principle
- Example:
 - An Image class holds image data and can read and write it from disk.
 - A second class, ImageFilter, has methods that manipulate Image objects and return new ones.

- **Resource Allocation Is Initialization (RAII):**

- A late 80’s concept, widely used in OOP.
- https://en.wikipedia.org/wiki/Resource_acquisition_is_initialization
- ALL Resources in a class are created in the constructor and released in the destructor.
 - Example: opening files, allocating memory, etc.
- If an object is created it is ready to use.

Further learning

- When looking for C++ tutorials and guides, look for ones that use at least the C++11 standard.
 - This is “modern C++”
- Some tutorials:
 - <https://cplusplus.com/doc/tutorial/>
 - <https://www.w3schools.com/CP/default.asp>
- Books:
 - [Effective Modern C++](#)
 - [The C++ Programming Language](#)