Design Patterns 1

6.170 Lecture 17

Spring 2004

1 Design patterns

A design pattern is:

- ▷ a standard solution to a common programming problem
- ▷ a technique for making code more flexible by making it meet certain criteria
- ▷ a design or implementation structure that achieves a particular purpose
- ▷ a high-level programming idiom
- ▷ shorthand for describing certain aspects of program organization
- ▷ connections among program components
- ▷ the shape of an object diagram or object model

1.1 Examples

Here are some examples of design patterns which you have already seen. For each design pattern, this list notes the problem it is trying to solve, the solution that the design pattern supplies, and any disadvantages associated with the design pattern. A software designer must trade off the advantages against the disadvantages when deciding whether to use a design pattern. Tradeoffs between flexibility and performance are common, as you will often discover in computer science (and other fields).

Encapsulation (data hiding)

Problem: Exposed fields can be directly manipulated from outside, leading to violations of the representation invariant or undesirable dependences that prevent changing the implementation.

Solution: Hide some components, permitting only stylized access to the object.

Disadvantages: The interface may not (efficiently) provide all desired operations. Indirection may reduce performance.

Subclassing (inheritance)

Problem: Similar abstractions have similar members (fields and methods). Repeating these is tedious, error-prone, and a maintenance headache.

Solution: Inherit default members from a superclass; select the correct implementation via run-time dispatching.

Disadvantages: Code for a class is spread out, potentially reducing understandability. Run-time dispatching introduces overhead.

Iteration

Problem: Clients that wish to access all members of a collection must perform a specialized traversal for each data structure. This introduces undesirable dependences and does not extend to other collections.

Solution: Implementations, which have knowledge of the representation, perform traversals and do bookkeeping. The results are communicated to clients via a standard interface.

Disadvantages: Iteration order is fixed by the implementation and not under the control of the client.

Exceptions

Problem: Errors occurring in one part of the code should often be handled elsewhere. Code should not be cluttered with error-handling code, nor return values preempted by error codes.

Solution: Introduce language structures for throwing and catching exceptions.

Disadvantages: Code may still be cluttered. It can be hard to know where an exception will be handled. Programmers may be tempted to use exceptions for normal control flow, which is confusing and usually inefficient.

These particular design patterns are so important that they are built into Java. Other design patterns are so important that they are built into other languages. Some design patterns may never be built into languages, but are still useful in their place.

1.2 When (not) to use design patterns

The first rule of design patterns is the same as the first rule of optimization: delay. Just as you shouldn't optimize prematurely, don't use design patterns prematurely. It may be best to first implement something and ensure that it works, then use the design pattern to improve weaknesses; this is especially true if you do not yet grasp all the details of the design. (If you fully understand the domain and problem, it may make sense to use design patterns from the start, just as it makes sense to use a more efficient rather than a less efficient algorithm from the very beginning in some applications.)

Design patterns may increase or decrease the understandability of a design or implementation. They can decrease understandability by adding indirection or increasing the amount of code. They can increase understandability by improving modularity, better separating concerns, and easing description. Once you learn the vocabulary of design patterns, you will be able to communicate more precisely and rapidly with other people who know the vocabulary. It's much better to say, "This is an instance of the visitor pattern" than "This is some code that traverses a structure and makes callbacks, and some certain methods must be present, and they are called in this particular way and in this particular order."

Most people use design patterns when they notice a problem with their design — something that ought to be easy isn't — or their implementation — such as performance. Examine the offending design or code. What are its problems, and what compromises does it make? What would you like to do that is presently too hard? Then, check a design pattern reference. Look for patterns that address the issues you are concerned with.

2 Creational patterns

2.1 Factories

Suppose you are writing a class to represent a bicycle race. A race consists of many bicycles (among other objects, perhaps).

```
class Race {

  Race createRace() {
    Frame frame1 = new Frame();
    Wheel frontWheel1 = new Wheel();
    Wheel rearWheel1 = new Wheel();
    Bicycle bike1 = new Bicycle(frame1, frontWheel1, rearWheel1);
    Frame frame2 = new Frame();
    Wheel frontWheel2 = new Wheel();
    Wheel rearWheel2 = new Wheel();
    Bicycle bike2 = new Bicycle(frame2, frontWheel2, rearWheel2);
    ...
  }

}
```

You can specialize `Race` for other bicycle races:

```
// French race
class TourDeFrance extends Race {

  Race createRace() {
    Frame frame1 = new RacingFrame();
    Wheel frontWheel1 = new Wheel700c();
    Wheel rearWheel1 = new Wheel700c();
    Bicycle bike1 = new Bicycle(frame1, frontWheel1, rearWheel1);
    Frame frame2 = new RacingFrame();
    Wheel frontWheel2 = new Wheel700c();
    Wheel rearWheel2 = new Wheel700c();
    Bicycle bike2 = new Bicycle(frame2, frontWheel2, rearWheel2);
    ...
  }

  ...
}

// all-terrain bicycle race
class Cyclocross extends Race {

  Race createRace() {
    Frame frame1 = new MountainFrame();
    Wheel frontWheel1 = new Wheel27in();
```

```
    Wheel rearWheel1 = new Wheel27in();
    Bicycle bike1 = new Bicycle(frame1, frontWheel1, rearWheel1);
    Frame frame2 = new MountainFrame();
    Wheel frontWheel2 = new Wheel27in();
    Wheel rearWheel2 = new Wheel27in();
    Bicycle bike2 = new Bicycle(frame2, frontWheel2, rearWheel2);
    ...
  }

  ...
}
```

In the subclasses, `createRace` returns a `Race` because the Java compiler enforces that overridden methods have identical return types.

For brevity, the code fragments above omit many other methods relating to bicycle races, some of which appear in each class and others of which appear only in certain classes.

The repeated code is tedious, and in particular, we weren't able to reuse method `Race.createRace` at all. (There is a separate issue of abstracting out the creation of a single bicycle to a function; we will use that without further discussion, as it is obvious, at least after 6.001.) There must be a better way. The Factory design patterns provide an answer.

2.1.1 Factory method

A factory method is a method that manufactures objects of a particular type.

We can add factory methods to `Race`:

```
class Race {

  Frame createFrame() { return new Frame(); }
  Wheel createWheel() { return new Wheel(); }
  Bicycle createBicycle(Frame frame, Wheel front, Wheel rear) {
    return new Bicycle(frame, front, rear);
  }
  // return a complete bicycle without needing any arguments
  Bicycle completeBicycle() {
    Frame frame = createFrame();
    Wheel frontWheel = createWheel();
    Wheel rearWheel = createWheel();
    return createBicycle(frame, frontWheel, rearWheel);
  }

  Race createRace() {
    Bicycle bike1 = completeBicycle();
    Bicycle bike2 = completeBicycle();
    ...
  }

}
```

Now subclasses can reuse `createRace` and even `completeBicycle` without change:

```
// French race
class TourDeFrance extends Race {

  Frame createFrame() { return new RacingFrame(); }
  Wheel createWheel() { return new Wheel700c(); }
  Bicycle createBicycle(Frame frame, Wheel front, Wheel rear) {
    return new RacingBicycle(frame, front, rear);
  }

}

class Cyclocross extends Race {

  Frame createFrame() { return new MountainFrame(); }
  Wheel createWheel() { return new Wheel26inch(); }
  Bicycle createBicycle(Frame frame, Wheel front, Wheel rear) {
    return new RacingBicycle(frame, front, rear);
  }

}
```

The `create...` methods are called *factory methods*.

2.1.2 Factory object

If there are many objects to construct, including the factory methods in each class can bloat the code and make it hard to change. Sibling subclasses cannot easily share the same factory method.

A *factory object* is an object that encapsulates factory methods.

```
class BicycleFactory {
  Frame createFrame() { return new Frame(); }
  Wheel createWheel() { return new Wheel(); }
  Bicycle createBicycle(Frame frame, Wheel front, Wheel rear) {
    return new Bicycle(frame, front, rear);
  }

  // return a complete bicycle without needing any arguments
  Bicycle completeBicycle() {
    Frame frame = createFrame();
    Wheel frontWheel = createWheel();
    Wheel rearWheel = createWheel();
    return createBicycle(frame, frontWheel, rearWheel);
  }
}

class RacingBicycleFactory {
  Frame createFrame() { return new RacingFrame(); }
```

```
  Wheel createWheel() { return new Wheel700c(); }
  Bicycle createBicycle(Frame frame, Wheel front, Wheel rear) {
    return new RacingBicycle(frame, front, rear);
  }
}

class MountainBicycleFactory {
  Frame createFrame() { return new MountainFrame(); }
  Wheel createWheel() { return new Wheel26inch(); }
  Bicycle createBicycle(Frame frame, Wheel front, Wheel rear) {
    return new RacingBicycle(frame, front, rear);
  }
}
```

The `Race` methods use the factory objects.

```
class Race {

  BicycleFactory bfactory;

  // constructor
  Race() {
    bfactory = new BicycleFactory();
  }

  Race createRace() {
    Bicycle bike1 = bfactory.completeBicycle();
    Bicycle bike2 = bfactory.completeBicycle();
    ...
  }

}

class TourDeFrance extends Race {
  // constructor
  TourDeFrance() {
    bfactory = new RacingBicycleFactory();
  }
}

class Cyclocross extends Race {
  // constructor
  Cyclocross() {
    bfactory = new MountainBicycleFactory();
  }
}
```

In this version of the code, the type of bicycle is still hard-coded into each variety of race. There is a more flexible method which requires a change to the way that clients call the constructor.

```
class Race {

  BicycleFactory bfactory;

  // constructor
  Race(BicycleFactory bfactory) {
    this.bfactory = bfactory;
  }

  Race createRace() {
    Bicycle bike1 = bfactory.completeBicycle();
    Bicycle bike2 = bfactory.completeBicycle();
    ...
  }

}

class TourDeFrance extends Race {
  // constructor
  TourDeFrance(BicycleFactory bfactory) {
    this.bfactory = bfactory;
  }
}

class Cyclocross extends Race {
  // constructor
  Cyclocross(BicycleFactory bfactory) {
    this.bfactory = bfactory;
  }
}
```

This is the most flexible mechanism of all. Now a client can control both the variety of race and the variety of bicycle used in the race, for instance via a call like

```
new TourDeFrance(new TricycleFactory())
```

One reason that factory methods are required is the *first weakness of Java constructors*: Java constructors always return an object of the specified type. They can never return an object of a subtype, even though that would be type-correct (both according to Java subtyping and according to true behavior subtyping as was described in previous lectures).

Factory methods can also overcome the *second weakness of Java constructors*: Java constructors always return a new object, and can never reuse existing objects. Factory methods don't have that limitation. Consider the following example:

```
class Boolean {
  ...
  // Convert primitive boolean to Boolean
  public static Boolean valueOf(boolean b) {
     return (b ? Boolean.TRUE : Boolean.FALSE);
```

```
        // Never needs to construct a new object
        // A constructor would be forced to do so
    }
    ...
}
```

2.2 Singleton

The singleton pattern guarantees that only one object of a particular class ever exists. This is useful for classes representing unique physical resources, such as the display. For example, to ensure that there can only ever be one Elvis, you could do this:

```
// Singleton with static factory
public class Elvis {
   private static final Elvis INSTANCE = new Elvis();
   private Elvis() { ... }
   public static Elvis getInstance() {
      return INSTANCE;
   }
   ...
}
```

The singleton pattern is also useful for large, expensive objects that should not be multiply instantiated. The reason that a factory method, rather than a constructor, must be used is the second weakness of Java constructors: Java constructors always return a new object, never a pre-existing object.

3 Structural patterns

Wrappers modify the behavior of another class; they are usually a thin veneer over the encapsulated class, which does the real work. The wrapper may modify the interface, extend the behavior, or restrict access. The wrapper intermediates between two incompatible interfaces, translating calls between the interfaces. This permits two pieces of code that were not designed or written together, and thus are slightly incompatible, to be used together anyway.

Three varieties of wrappers are adapters, decorators, and proxies:

Pattern	Functionality	Interface
Adapter	same	different
Decorator	different	same
Proxy	same	same

The functionality and interfaces compared are those at the inside and outside of the wrapper; that is, a client's view of the wrapped object is compared to a client's view of the wrapper.

3.1 Adapter

Adapters change the interface of a class without changing its basic functionality. For instance, they might permit interoperability between a geometry package that requires angles to be specified in radians and a client that expects to pass angles in degrees. Here are two other examples:

3.1.1 Example: Rectangle

Suppose that you have written code that works on `Rectangle` objects and calls their `scale` method.

```
interface Rectangle {
  // grow or shrink this by the given factor
  void scale(float factor);

  // other operations
  float area();
  float circumference();
  ...
}

class myClass {

  void myMethod(Rectangle r) {
    ...
    r.scale(2);
    ...
  }

}
```

Suppose there is another class `NonScaleableRectangle` which lacks the `scale` method but does have the other methods of `Rectangle`, as well as additional `setWidth` and `setHeight` methods.

```
class NonScaleableRectangle {
  void setWidth(float width) { ... }
  void setHeight(float height) { ... }
  ...
}
```

You may wish to switch to (or at least permit use of) this variety of rectangle, perhaps because it has desirable features, such as better performance, or perhaps because it is used elsewhere, in a system with which you need to interoperate.

You cannot use `NonScaleableRectangle` directly because of the incompatible interface. However, you can write an adapter which permits its use. There are two ways to do this: subclassing and composition/forwarding. The subclassing solution will be familiar:

```
class ScaleableRectangle1 extends NonScaleableRectangle implements Rectangle {
  void scale(float factor) {
    setWidth(factor * getWidth());
    setHeight(factor * getHeight());
  }
}
```

Composition/forwarding is a technique for "passing the buck", forwarding a request so that a different object does the requested work.

```
class ScaleableRectangle2 implements Rectangle {
  NonScaleableRectangle r;
  ScaleableRectangle2(NonScaleableRectangle r) {
    this.r = r;
  }

  void scale(float factor) {
    setWidth(factor * getWidth());
    setHeight(factor * getHeight());
  }

  float area() { return r.area(); }
  float circumference() { return r.circumference(); }
  ...
}
```

3.1.2 Example: Palette

Suppose that Professor Guttag calls Professor Devadas late at night because someone has discovered a problem with the problem set: it needs to support bicycles that can be repainted (to change their color). The professors split up the work: Professor Guttag will write a `ColorPalette` class with a method that, given a name like “red” or “blue” or “taupe”, returns an array of three RGB values, and Professor Devadas will write code that uses this class. The professors do so, test their work, and go away for the weekend, leaving the the `.class` files for the TAs to integrate. They find that Professor Devadas has written code that depends on

```
interface ColorPalette {
  // returns RGB values
  int[] getColorPalette(String name);
}
```

but Professor Guttag has implemented a class that adheres to

```
interface ColorPallet {
  // returns RGB values
  int[] getColorPallet(String name);
}
```

What are the TAs to do? They do not have access to the source, and they do not have time to reimplement and retest. Their solution is to write an adapter for `ColorPallet` that changes the operation name. They can implement the adapter either by subclassing or by composition/forwarding, as described in Lecture 14.

3.2 Decorator

Whereas an adapter changes an interface without adding new functionality, a decorator extends functionality while maintaining the same interface. Typically, a decorator does not change existing functionality, only adds to it, so that objects of the resulting class behave exactly like the original ones, but also do something extra.

This sounds like subclassing, but not every instance of subclassing is a decoration. First, the implementation of an operation may be completely different or reimplemented in a subclass; that is not usually the case for a decorator, which contains relatively less functionality and reuses the superclass code. Second, subclasses can introduce new operations; wrappers (including decorators) generally do not.

An example of decoration is a `Window` interface (for a window manager) and a `BorderedWindow` interface. The `BorderedWindow` behaves exactly like the `Window`, except that it also draws a border around the outside.

Suppose that `Window` is implemented like this:

```
interface Window {
  // rectangle bounding the window
  Rectangle bounds();
  // draw this on the specified screen
  void draw(Screen s);
  ...
}

class WindowImpl implements Window {
  ...
}
```

The subclassing implementation would look like this:

```
class BorderedWindow1 extends WindowImpl {
  void draw(Screen s) {
    super.draw(s);
    bounds().draw(s);
  }
}
```

The composition/forwarding implementation would look like this:

```
class BorderedWindow2 implements Window {
  Window innerWindow;

  BorderedWindow2(Window innerWindow) {
    this.innerWindow = innerWindow;
  }

  void draw(Screen s) {
    innerWindow.draw(s);
    innerWindow.bounds().draw(s);
  }
}
```

3.3 Proxy

A proxy is a wrapper that has the same interface and the same functionality as the class it wraps. This does not sound very useful on the face of it. However, proxies serve an important purpose in

controlling access to other objects. This is particularly valuable if those objects must be accessed in a stylized or complicated way.

For example, if an object is on a remote machine, then accessing it requires use of various network facilities. It is easier to create a local proxy that understands the network and performs the necessary operations, then returns the result. This simplifies the client by localizing network-specific code in another location.

As another example, an object may require locking if it can be accessed by multiple clients. The lock represents the right to read and/or update the object; without the lock, concurrent updates could leave the object in an inconsistent state, or reads in the middle of a sequence of updates could observe an inconsistent state. A proxy could take care of locking an object before an operation or sequence of operations, then unlocking it afterward. This is less error-prone than requiring clients to correctly implement the locking protocol.

Another variety of proxy is a security proxy. It might operate correctly if the caller has the correct credentials (such as a valid Kerberos certificate), but throw an error if an unauthorized user attempts to perform operations.

A final example is a proxy for an object that may not yet exist. If creating an object is expensive (because of computation or network latency), then it can be represented by a proxy instead. That proxy could immediately start to create the object in a background task in the hope that it is ready by the time the first operation is invoked, or it could delay creating the object until an operation is invoked. In the former case, the rest of the system can proceed without waiting; in the latter case, the work of creating the object need never be performed if it is never used. In either case, operations are delayed until the object is ready.

An example of a proxy for a non-existent object is Emacs's autoload functionality. For instance, suppose I have a file `util-mde.el` which defines of a number of useful functions. However, I don't want to slow down Emacs by loading it every time I start Emacs. Instead, my `.emacs` file contains code like this:

```
(autoload 'looking-back-at "util-mde")
(autoload 'in-buffer "util-mde")
(autoload 'in-window "util-mde")
```

The form `(autoload 'function "file")` is essentially equivalent to (in Scheme syntax; Emacs Lisp uses `defun`)

```
(define function ()
  (load "file")  ;; redefine function
  (function)     ;; call the new version
)
```

Emacs autoloads most of its own functionality, from the mail and news readers to the Java editing mode. People who complain that Emacs starts up too slowly often have put indiscriminate `load` forms in their `.emacs` files; that's like using an inefficient implementation, then complaining that the compiler is poor because the resulting program runs slowly.

Proxy capabilities are particularly useful when clients have no knowledge of whether the object they are manipulating has special properties (such as being located on a remote machine, requiring locking or security, or not being loaded). It is best to insulate the client from such concerns and localize them in a proxy wrapper.

3.4 Composite

The composite design pattern permits a client to manipulate either an atomic unit or a collection of units in exactly the same way. The client need not create special code for the case of being given a higher-level object with structure as opposed to being given a basic object; the same operations work on both. Composite is good for object with part-whole relationships, and the client should not have to worry about whether its argument is atomic or composed of parts. For example, a bicycle can be decomposed in the following way:

```
Bicycle
  Wheel
    skewer
    hub
    spokes
    nipples
    rim
    tube
    tire
  Frame
  Drivetrain
  ...
```

Given a bicycle component, I might want to determine its weight or cost regardless of whether it is itself composed of subcomponents. A client which is given a bicycle component shouldn't have to treat it differently if it is a wheel as opposed to a reflector or a saddle.

The solution to this problem is for all bicycle components to satisfy a common interface:

```
class BicycleComponent {
  int weight();
  float cost();
}
```

The implementation of `Wheel.weight` might itself call `weight` on its subparts, but that is of no import to the client (and the client shouldn't have to worry about that). An alternative to using a common interface is to have a common superclass; in either way, all bicycle components at all levels provide the same methods and can be used interchangeably.

design patterns

Standard solutions to common programming problems

Distilled, debugged, documented, sometimes over decades

We've been using some of them all term
- encapsulation, inheritance, exception, iterator, etc.

Lots of info available about trade-offs and variants

Good to know about, but don't apply them indiscriminately

when (not) to use design patterns

Delay – get something basic working first, then improve it once you understand it

Design patterns can increase or decrease understandability
- Add indirection, increase code size
- Improve modularity, separate concerns, ease description

If your design has a problem, consider design patterns that address that problem

Canonical reference: the "Gang of Four" book
Good Java reference: the Bloch book

categories of patterns

Creational patterns
- Minimize coupling introduced by object creation

Structural patterns
- Minimize coupling introduced when classes or objects are assembled to form larger structures

Behavioral patterns
- Minimize coupling introduced when classes or objects cooperate to perform distributed tasks

categories of patterns

Creational patterns
- Minimize coupling introduced by object creation

Structural patterns
- Minimize coupling introduced when classes or objects are assembled to form larger structures

Behavioral patterns
- Minimize coupling introduced when classes or objects cooperate to perform distributed tasks

problem

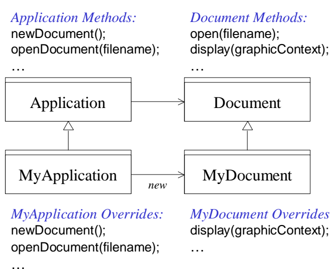

Application+Document classes handle UI for opening/editing multiple documents. How can we allow customization of documents?

solution: Factory Method pattern

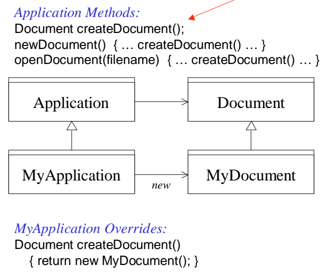

Give Application a *createDocument()* method which can be overridden appropriately. Never call *new Document()*, always *createDocument()*

advantages of Factory Method

Hides decision about what concrete class to create

Can overcome two limitations of Java constructors

1. Constructor can only return an instance of its associated class, never a subtype
2. Constructor must always return a new object, never a previously allocated object

reusing objects with Factory Method

```
class Boolean {
  …
  // Convert primitive boolean to Boolean.
  public static Boolean valueOf(boolean b) {
    return (b ? Boolean.TRUE : Boolean.FALSE);
    // Never needs to construct a new object.
    // A constructor would be forced to do so.
  }
  …
}
```

problem

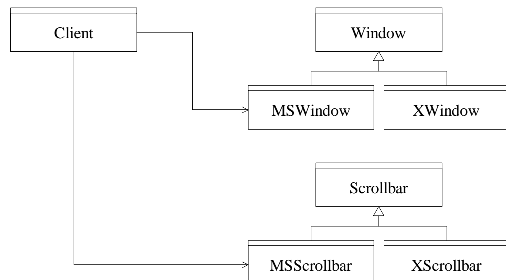

Want to insulate client code from details of different windowing platforms or look-and-feel standards

solution: Abstract Factory pattern

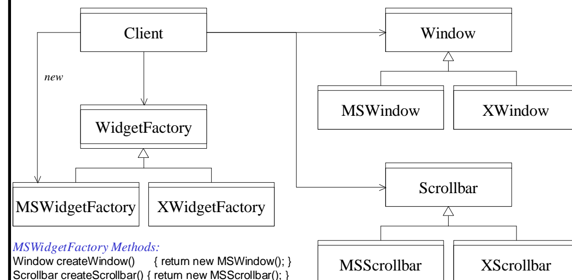

Make a "factory" class that knows how to construct the local kind of widgets; client creates all graphical objects using factory

problem

```
// the one and only
public class Elvis {
  …
}
```

Want to guarantee that there is exactly one Elvis object in system. Useful for e.g. scheduler, file system, video display driver, etc.

solution: Singleton pattern

```
// Singleton with static factory
public class Elvis {
  private static final Elvis INSTANCE = new Elvis();
  private Elvis() { … }
  public static Elvis getInstance() {
    return INSTANCE;
  }
  …
}
```

There are several ways to do this, depending on whether you might eventually need more than one Elvis (scheduler, file system, …)

categories of patterns

Creational patterns

Minimize coupling introduced by object creation

Structural patterns

Minimize coupling introduced when classes or objects are assembled to form larger structures

Behavioral patterns

Minimize coupling introduced when classes or objects cooperate to perform distributed tasks

problem

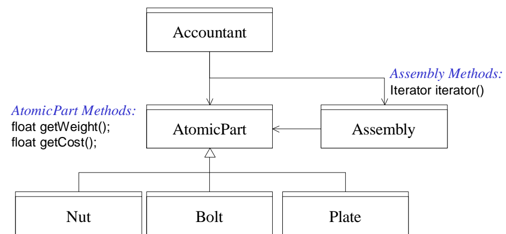

Parts can be atomic or assembled out of other parts. Want to let client access total cost/weight of parts without distinguishing atomic/assembly.

solution: Composite pattern

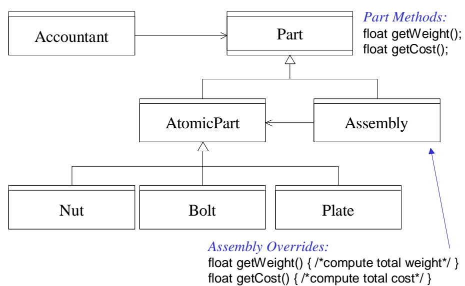

Implement a uniform interface, with assemblies aggregating information across their content instead of the client (e.g. totaling cost, weight)

problem

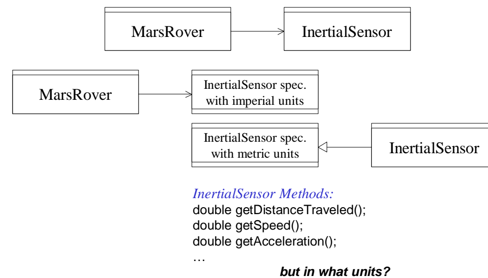

Through an unfortunate misunderstanding, the InertialSensor is implemented using metric units, not the expected imperial units

solution: Adapter pattern

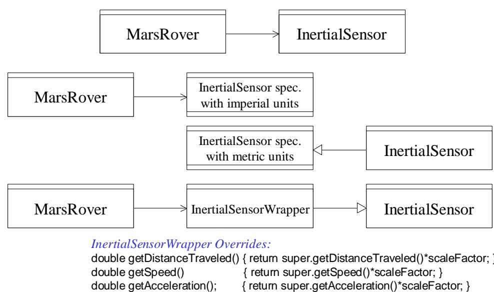

Of course NASA catches the problem early in testing, and fixes it with a Adapter, using subclassing or composition/forwarding to fix interface

another example for Adaptor

```
// GUI expects rectangle with scale() method
interface ScaleRectangle {
  // grow or shrink by the given factor
  void scale(float factor);
  float getCircumference();
  …
}

// but instead we have this
class DimensionRectangle {
  void setWidth(float width) { … }
  void setHeight(float height) { … }
  float getCircumference() { … }
  …
}
```

```
// apply Adapter pattern
class PseudoScaleRectangle
    extends DimensionRectangle
    implements ScaleRectangle {
  void scale(float factor) {
    setWidth(factor*getWidth());
    setHeight(factor*getHeight());
  }
}
```

another example for Adaptor

```
// apply Adapter pattern using composition/forwarding
class PseudoScaleRectangle implements ScaleRectangle {
  DimensionRectangle r;
  PseudoScaleRectangle(DimensionRectangle r) {
    this.r = r;
  }
  void scale(float factor) {
    r.setWidth(factor*r.getWidth());
    r.setHeight(factor*r.getHeight());
  }
  float getCircumference() { return r.getCircumference(); }
  …
}
```

Composition/forwarding approach completely hides DimensionRectangle from client

Wrapper patterns

Adapter

- Same functionality, different interface
  - Rename a method, convert units, implement one method in terms of another, etc.

Decorator

- Different functionality, same interface
  - Change methods to do more (while still meeting original spec)

Proxy

- Same functionality, same interface
  - Represent remote objects, add security, allow lazy creation, etc.

example for Decorator

```
interface Window {
  // rectangle bounding the window
  Rectangle bounds();
  // draw this window in the specified part of the user interface
  void draw(GraphicContext g);
  …
}

class WindowImpl implements Window {
 …
}
```

Suppose we want to add borders to our windows, but can't modify WindowImpl code

example for Decorator

```
class BorderedWindow1 extends WindowImpl {
  void draw(GraphicContext g) {
    super.draw(g);
    bounds().draw(g);
  }
}
```

Subclassing approach

```
class BorderedWindow2 implements Window {
  Window innerWindow;
  BorderedWindow2(Window w) { this.innerWindow = w; }
  void draw(GraphicContext g) {
    innerWindow.draw(g);
    bounds().draw(g);
  }
}
```

Composition/forwarding approach

example for Proxy

```
class EmbeddedImage {
  EmbeddedImage(String filename)  {
    // loads image from file
  }
  void draw(GraphicContext g)  {
    // renders image
  }
  …
}
```

EmbeddedImage represents images placed into a document. All images are loaded when the document is first opened, giving sluggish behavior. Can we delay loading?

example for Proxy

```
class EmbeddedImageProxy {
  EmbeddedImage image;
  String filename;
  EmbeddedImageProxy(String filename)  {
    this.filename = filename;
  }
 void draw(GraphicContext g) {
    if (image==null)  {
      image = new EmbeddedImage(filename);
    }
    image.draw(g);
  }
  …
}
```