



Some Pattern Principles

Although design patterns simply describe solutions to problems, they tend to emphasize solutions that promote reusability and flexibility. To achieve this, they manifest some key object-oriented design principles. We will encounter some of them in this chapter and in more detail throughout the rest of the book.

This chapter will cover

- *Composition*: How to use object aggregation to achieve greater flexibility than you could with inheritance alone
- *Decoupling*: How to reduce dependency between elements in a system
- *The power of the interface*: Patterns and polymorphism
- *Pattern categories*: The types of pattern that this book will cover

The Pattern Revelation

I first started working with objects in the Java language. As you might expect, it took a while before some concepts clicked. When it did happen, though, it happened very fast, almost with the force of revelation. The elegance of inheritance and encapsulation bowled me over. I could sense that this was a different way of defining and building systems. I *got* polymorphism, working with a type and switching implementations at runtime.

All the books on my desk at the time focused on language features and the very many APIs available to the Java programmer. Beyond a brief definition of polymorphism, there was little attempt to examine design strategies.

Language features alone do not engender object-oriented design. Although my projects fulfilled their functional requirements, the kind of design that inheritance, encapsulation, and polymorphism had seemed to offer continued to elude me.

My inheritance hierarchies grew wider and deeper as I attempted to build new classes for every eventuality. The structure of my systems made it hard to convey messages from one tier to another without giving intermediate classes too much awareness of their surroundings, binding them into the application and making them unusable in new contexts.

It wasn't until I discovered *Design Patterns*, otherwise known as the Gang of Four book, that I realized I had missed an entire design dimension. By that time, I had already discovered some of the core patterns for myself, but others contributed to a new way of thinking.

I discovered that I had overprivileged inheritance in my designs, trying to build too much functionality into my classes. But where else can functionality go in an object-oriented system?

I found the answer in composition. Software components can be defined at runtime by combining objects in flexible relationships. The Gang of Four boiled this down into a principle: “favor composition

over inheritance.” The patterns described ways in which objects could be combined at runtime to achieve a level of flexibility impossible in an inheritance tree alone.

Composition and Inheritance

Inheritance is a powerful way of designing for changing circumstances or contexts. It can limit flexibility, however, especially when classes take on multiple responsibilities.

The Problem

As you know, child classes inherit the methods and properties of their parents (as long as they are protected or public elements). You can use this fact to design child classes that provide specialized functionality.

Figure 8–1 presents a simple example using the UML.

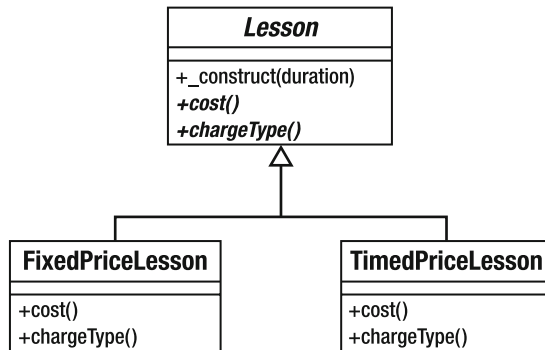


Figure 8–1. A parent class and two child classes

The abstract `Lesson` class in Figure 8–1 models a lesson in a college. It defines abstract `cost()` and `chargeType()` methods. The diagram shows two implementing classes, `FixedPriceLesson` and `TimedPriceLesson`, which provide distinct charging mechanisms for lessons.

Using this inheritance scheme, I can switch between lesson implementations. Client code will know only that it is dealing with a `Lesson` object, so the details of `cost` will be transparent.

What happens, though, if I introduce a new set of specializations? I need to handle lectures and seminars. Because these organize enrollment and lesson notes in different ways, they require separate classes. So now I have two forces that operate upon my design. I need to handle pricing strategies and separate lectures and seminars.

Figure 8–2 shows a brute-force solution.




Figure 8–2. A poor inheritance structure

Figure 8–2 shows a hierarchy that is clearly faulty. I can no longer use the inheritance tree to manage my pricing mechanisms without duplicating great swathes of functionality. The pricing strategies are mirrored across the Lecture and Seminar class families.

At this stage, I might consider using conditional statements in the Lesson super class, removing those unfortunate duplications. Essentially, I remove the pricing logic from the inheritance tree altogether, moving it up into the super class. This is the reverse of the usual refactoring where you replace a conditional with polymorphism. Here is an amended Lesson class:

```

abstract class Lesson {
    protected $duration;
    const    FIXED = 1;
    const    TIMED = 2;
    private  $costtype;

    function __construct( $duration, $costtype=1 ) {
        $this->duration = $duration;
        $this->costtype = $costtype;
    }

    function cost() {
        switch ( $this->costtype ) {
            CASE self::TIMED :
                return ( 5 * $this->duration );
                break;
            CASE self::FIXED :
                return 30;
                break;
            default:
                $this->costtype = self::FIXED;
                return 30;
        }
    }
}
  
```

```

function chargeType() {
  switch ( $this->costtype ) {
    CASE self::TIMED :
      return "hourly rate";
      break;
    CASE self::FIXED :
      return "fixed rate";
      break;
    default:
      $this->costtype = self::FIXED;
      return "fixed rate";
  }
}

// more lesson methods...
}

class Lecture extends Lesson {
  // Lecture-specific implementations ...
}

class Seminar extends Lesson {
  // Seminar-specific implementations ...
}

```

Here's how I might work with these classes:

```

$lecture = new Lecture( 5, Lesson::FIXED );
print "{$lecture->cost()} ({$lecture->chargeType()})\n";

$seminar= new Seminar( 3, Lesson::TIMED );
print "{$seminar->cost()} ({$seminar->chargeType()})\n";

```

And here's the output:

```

30 (fixed rate)
15 (hourly rate)

```

You can see the new class diagram in Figure 8-3.

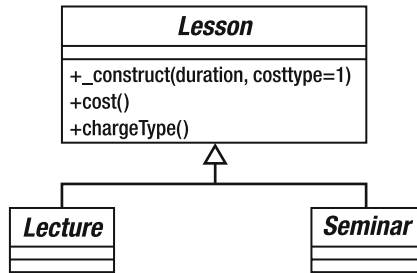


Figure 8–3. Inheritance hierarchy improved by removing cost calculations from subclasses

I have made the class structure much more manageable but at a cost. Using conditionals in this code is a retrograde step. Usually, you would try to replace a conditional statement with polymorphism. Here, I have done the opposite. As you can see, this has forced me to duplicate the conditional statement across the `chargeType()` and `cost()` methods.

I seem doomed to duplicate code.

Using Composition

I can use the Strategy pattern to compose my way out of trouble. Strategy is used to move a set of algorithms into a separate type. By moving cost calculations, I can simplify the Lesson type. You can see this in Figure 8–4.

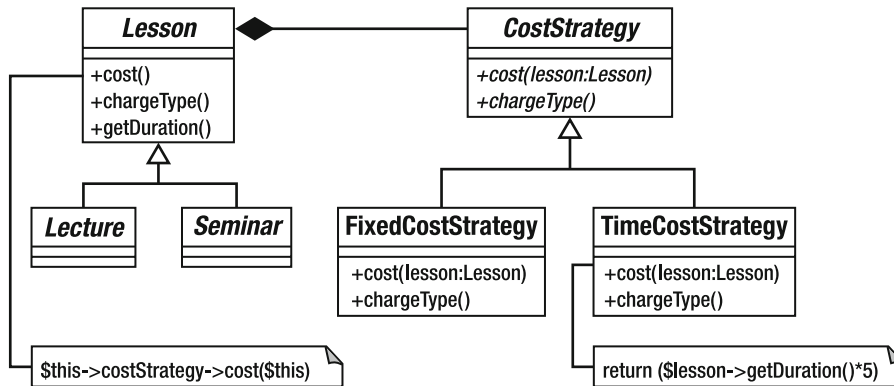


Figure 8–4. Moving algorithms into a separate type

I create an abstract class, `CostStrategy`, which defines the abstract methods `cost()` and `chargeType()`. The `cost()` method requires an instance of `Lesson`, which it will use to generate cost data. I provide two implementations for `CostStrategy`. `Lesson` objects work only with the `CostStrategy` type, not a specific implementation, so I can add new cost algorithms at any time by subclassing `CostStrategy`. This would require no changes at all to any `Lesson` classes.

Here's a simplified version of the new `Lesson` class illustrated in Figure 8–4:

```

abstract class Lesson {
  private $duration;
  private $costStrategy;

  function __construct( $duration, CostStrategy $strategy ) {
    $this->duration = $duration;
    $this->costStrategy = $strategy;
  }

  function cost() {
    return $this->costStrategy->cost( $this );
  }

  function chargeType() {
    return $this->costStrategy->chargeType( );
  }

  function getDuration() {
    return $this->duration;
  }

  // more lesson methods...
}

class Lecture extends Lesson {
  // Lecture-specific implementations ...
}

class Seminar extends Lesson {
  // Seminar-specific implementations ...
}

```

The Lesson class requires a CostStrategy object, which it stores as a property. The Lesson::cost() method simply invokes CostStrategy::cost(). Equally, Lesson::chargeType() invokes CostStrategy::chargeType(). This explicit invocation of another object's method in order to fulfill a request is known as delegation. In my example, the CostStrategy object is the delegate of Lesson. The Lesson class washes its hands of responsibility for cost calculations and passes on the task to a CostStrategy implementation. Here, it is caught in the act of delegation:

```

function cost() {
  return $this->costStrategy->cost( $this );
}

```

Here is the CostStrategy class, together with its implementing children:

```

abstract class CostStrategy {
  abstract function cost( Lesson $lesson );
  abstract function chargeType();
}

class TimedCostStrategy extends CostStrategy {
  function cost( Lesson $lesson ) {
    return ( $lesson->getDuration() * 5 );
  }
  function chargeType() {

```

```

        return "hourly rate";
    }
}

class FixedCostStrategy extends CostStrategy {
    function cost( Lesson $lesson ) {
        return 30;
    }

    function chargeType() {
        return "fixed rate";
    }
}

```

I can change the way that any Lesson object calculates cost by passing it a different CostStrategy object at runtime. This approach then makes for highly flexible code. Rather than building functionality into my code structures statically, I can combine and recombine objects dynamically.

```

$lessons[] = new Seminar( 4, new TimedCostStrategy() );
$lessons[] = new Lecture( 4, new FixedCostStrategy() );

foreach ( $lessons as $lesson ) {
    print "lesson charge {$lesson->cost()}. ";
    print "Charge type: {$lesson->chargeType()}\n";
}

```

lesson charge 20. Charge type: hourly rate

lesson charge 30. Charge type: fixed rate

As you can see, one effect of this structure is that I have focused the responsibilities of my classes. CostStrategy objects are responsible solely for calculating cost, and Lesson objects manage lesson data.

So, composition can make your code more flexible, because objects can be combined to handle tasks dynamically in many more ways than you can anticipate in an inheritance hierarchy alone. There can be a penalty with regard to readability, though. Because composition tends to result in more types, with relationships that aren't fixed with the same predictability as they are in inheritance relationships, it can be slightly harder to digest the relationships in a system.

Decoupling

You saw in Chapter 6 that it makes sense to build independent components. A system with highly interdependent classes can be hard to maintain. A change in one location can require a cascade of related changes across the system.

The Problem

Reusability is one of the key objectives of object-oriented design, and tight coupling is its enemy. You can diagnose tight coupling when you see that a change to one component of a system necessitates many changes elsewhere. You should aspire to create independent components so that you can make changes without a domino effect of unintended consequences. When you alter a component, the extent

to which it is independent is related to the likelihood that your changes will cause other parts of your system to fail.

You saw an example of tight coupling in Figure 8–2. Because the costing logic was mirrored across the `Lecture` and `Seminar` types, a change to `TimedPriceLecture` would necessitate a parallel change to the same logic in `TimedPriceSeminar`. By updating one class and not the other, I would break my system—without any warning from the PHP engine. My first solution, using a conditional statement, produced a similar dependency between the `cost()` and `chargeType()` methods.

By applying the Strategy pattern, I distilled my costing algorithms into the `CostStrategy` type, locating them behind a common interface and implementing each only once.

Coupling of another sort can occur when many classes in a system are embedded explicitly into a platform or environment. Let's say that you are building a system that works with a MySQL database, for example. You might use functions such as `mysql_connect()` and `mysql_query()` to speak to the database server.

Should you be required to deploy the system on a server that does not support MySQL, you *could* convert your entire project to use SQLite. You would be forced to make changes throughout your code, though, and face the prospect of maintaining two parallel versions of your application.

The problem here is not the system's dependency on an external platform. Such a dependency is inevitable. You need to work with code that speaks to a database. The problem comes when such code is scattered throughout a project. Talking to databases is not the primary responsibility of most classes in a system, so the best strategy is to extract such code and group it together behind a common interface. In this way, you promote the independence of your classes. At the same time, by concentrating your gateway code in one place, you make it much easier to switch to a new platform without disturbing your wider system. This process, the hiding of implementation behind a clean interface, is known as encapsulation.

PEAR solves this problem with the `PEAR::MDB2` package (which has superseded `PEAR::DB`). This provides a single point of access for multiple databases. More recently the bundled PDO extension has brought this model into the PHP language itself.

The `MDB2` class provides a static method called `connect()` that accepts a Data Source Name (DSN) string. According to the makeup of this string, it returns a particular implementation of a class called `MDB2_Driver_Common`. So for the string `"mysql://"`, the `connect()` method returns a `MDB2_Driver_mysql` object, while for a string that starts with `"sqlite://"`, it would return an `MDB2_Driver_sqlite` object. You can see the class structure in Figure 8–5.




Figure 8–5. The `PEAR::MDB2` package decouples client code from database objects.

The `PEAR::MDB2` package, then, lets you decouple your application code from the specifics of your database platform. As long as you use uncontroversial SQL, you should be able to run a single system with MySQL, SQLite, MSSQL, and others without changing a line of code (apart from the DSN, of course, which is the single point at which the database context must be configured). In fact, the `PEAR::MDB2` package can also help manage different SQL dialects to some extent—one reason you might still choose to use it, despite the speed and convenience of PDO.

Loosening Your Coupling

To handle database code flexibly, you should decouple the application logic from the specifics of the database platform it uses. You will see lots of opportunities for this kind of component separation of components in your own projects.

Imagine for example that the Lesson system must incorporate a registration component to add new lessons to the system. As part of the registration procedure, an administrator should be notified when a lesson is added. The system's users can't agree whether this notification should be sent by mail, or by text message. In fact, they're so argumentative, that you suspect they might want to switch to a new mode of communication in the future. What's more, they want to be notified of all sorts of things. So that a change to the notification mode in one place, will mean a similar alteration in many other places.

If you've hardcoded calls to a Mailer class, or a Texter class, then your system is tightly coupled to a particular notification mode. Just as it would be tightly coupled to a database platform by the use of a specialized database API.

Here is some code that hides the implementation details of a notifier from the system that uses it.

```
class RegistrationMgr {
    function register( Lesson $lesson ) {
        // do something with this Lesson

        // now tell someone
        $notifier = Notifier::getNotifier();
        $notifier->inform( "new lesson: cost ({$lesson->cost()})" );
    }
}
```

```
abstract class Notifier {

    static function getNotifier() {
        // acquire concrete class according to
        // configuration or other logic

        if ( rand(1,2) == 1 ) {
            return new MailNotifier();
        } else {
```

```

        return new TextNotifier();
    }
}

abstract function inform( $message );
}

class MailNotifier extends Notifier {
    function inform( $message ) {
        print "MAIL notification: {$message}\n";
    }
}

class TextNotifier extends Notifier {
    function inform( $message ) {
        print "TEXT notification: {$message}\n";
    }
}

```

I create `RegistrationMgr`, a sample client for my `Notifier` classes. The `Notifier` class is abstract, but it does implement a static method: `getNotifier()` which fetches a concrete `Notifier` object (`TextNotifier` or `MailNotifier`). In a real project, the choice of `Notifier` would be determined by a flexible mechanism, such as a configuration file. Here, I cheat and make the choice randomly. `MailNotifier` and `TextNotifier` do nothing more than print out the message they are passed along with an identifier to show which one has been called.

Notice how the knowledge of which concrete `Notifier` should be used has been focused in the `Notifier::getNotifier()` method. I could send notifier messages from a hundred different parts of my system, and a change in `Notifier` would only have to be made in that one method.

Here is some code that calls the `RegistrationMgr`,

```

$lessons1 = new Seminar( 4, new TimedCostStrategy() );
$lessons2 = new Lecture( 4, new FixedCostStrategy() );
$mgr = new RegistrationMgr();

```

```
$mgr->register( $lessons1 );
$mgr->register( $lessons2 );
```

and the output from a typical run

```
TEXT notification: new lesson: cost (20)
MAIL notification: new lesson: cost (30)
```

Figure 8–6 shows these classes.




Figure 8–6. The *Notifier* class separates client code from *Notifier* implementations.

Notice how similar the structure in Figure 8–6 is to that formed by the MDB2 components shown in Figure 8–5

Code to an Interface, Not to an Implementation

This principle is one of the all-pervading themes of this book. You saw in Chapter 6 (and in the last section) that you can hide different implementations behind the common interface defined in a superclass. Client code can then require an object of the superclass’s type rather than that of an implementing class, unconcerned by the specific implementation it is actually getting.

Parallel conditional statements, like the ones I built into `Lesson::cost()` and `Lesson::chargeType()`, are a common signal that polymorphism is needed. They make code hard to maintain, because a change in one conditional expression necessitates a change in its twins. Conditional statements are occasionally said to implement a “simulated inheritance.”

By placing the cost algorithms in separate classes that implement `CostStrategy`, I remove duplication. I also make it much easier should I need to add new cost strategies in the future.

From the perspective of client code, it is often a good idea to require abstract or general types in your methods’ parameters. By requiring more specific types, you could limit the flexibility of your code at runtime.

Having said that, of course, the level of generality you choose in your argument hints is a matter of judgment. Make your choice too general, and your method may become less safe. If you require the specific functionality of a subtype, then accepting a differently equipped sibling into a method could be risky.

Still, make your choice of argument hint too restricted, and you lose the benefits of polymorphism. Take a look at this altered extract from the `Lesson` class:

```
function __construct( $duration,
    FixedPriceStrategy $strategy ) {
    $this->duration = $duration;
    $this->costStrategy = $strategy;
}
```

There are two issues arising from the design decision in this example. First, the Lesson object is now tied to a specific cost strategy, which closes down my ability to compose dynamic components. Second, the explicit reference to the FixedPriceStrategy class forces me to maintain that particular implementation.

By requiring a common interface, I can combine a Lesson object with any CostStrategy implementation:

```
function __construct( $duration, CostStrategy $strategy ) {
    $this->duration = $duration;
    $this->costStrategy = $strategy;
}
```

I have, in other words, decoupled my Lesson class from the specifics of cost calculation. All that matters is the interface and the guarantee that the provided object will honor it.

Of course, coding to an interface can often simply defer the question of how to instantiate your objects. When I say that a Lesson object can be combined with any CostStrategy interface at runtime, I beg the question, “But where does the CostStrategy object come from?”

When you create an abstract super class, there is always the issue as to how its children should be instantiated. Which child do you choose and according to which condition? This subject forms a category of its own in the Gang of Four pattern catalog, and I will examine it further in the next chapter.

The Concept That Varies

It’s easy to interpret a design decision once it has been made, but how do you decide where to start?

The Gang of Four recommend that you “encapsulate the concept that varies.” In terms of my lesson example, the varying concept is the cost algorithm. Not only is the cost calculation one of two possible strategies in the example, but it is obviously a candidate for expansion: special offers, overseas student rates, introductory discounts, all sorts of possibilities present themselves.

I quickly established that subclassing for this variation was inappropriate, and I resorted to a conditional statement. By bringing my variation into the same class, I underlined its suitability for encapsulation.

The Gang of Four recommend that you actively seek varying elements in your classes and assess their suitability for encapsulation in a new type. Each alternative in a suspect conditional may be extracted to form a class extending a common abstract parent. This new type can then be used by the class or classes from which it was extracted. This has the effect of

- Focusing responsibility
- Promoting flexibility through composition
- Making inheritance hierarchies more compact and focused
- Reducing duplication

So how do you spot variation? One sign is the misuse of inheritance. This might include inheritance deployed according to multiple forces at one time (lecture/seminar, fixed/timed cost). It might also include subclassing on an algorithm where the algorithm is incidental to the core responsibility of the type. The other sign of variation suitable for encapsulation is, of course, a conditional expression.

Patternitis

One problem for which there is no pattern is the unnecessary or inappropriate use of patterns. This has earned patterns a bad name in some quarters. Because pattern solutions are neat, it is tempting to apply them wherever you see a fit, whether they truly fulfill a need or not.

The eXtreme Programming (XP) methodology offers a couple of principles that might apply here. The first is “You aren’t going to need it” (often abbreviated to YAGNI). This is generally applied to application features, but it also makes sense for patterns.

When I build large environments in PHP, I tend to split my application into layers, separating application logic from presentation and persistence layers. I use all sorts of core and enterprise patterns in conjunction with one another.

When I am asked to build a feedback form for a small business web site, however, I may simply use procedural code in a single page script. I do not need enormous amounts of flexibility, I won’t be building on the initial release. I don’t need to use patterns that address problems in larger systems. Instead, I apply the second XP principle: “Do the simplest thing that works.”

When you work with a pattern catalog, the structure and process of the solution are what stick in the mind, consolidated by the code example. Before applying a pattern, though, pay close attention to the problem, or “when to use it,” section, and read up on the pattern’s consequences. In some contexts, the cure may be worse than the disease.

The Patterns

This book is not a pattern catalog. Nevertheless, in the coming chapters, I will introduce a few of the key patterns in use at the moment, providing PHP implementations and discussing them in the broad context of PHP programming.

The patterns described will be drawn from key catalogs including *Design Patterns*, *Patterns of Enterprise Application Architecture* by Martin Fowler (Addison-Wesley, 2003) and *Core J2EE Patterns* by Alur et al. (Prentice Hall PTR, 2001). I use the Gang of Four’s categorization as a starting point, dividing patterns as follows.

Patterns for Generating Objects

These patterns are concerned with the instantiation of objects. This is an important category given the principle “code to an interface.” If you are working with abstract parent classes in your design, then you must develop strategies for instantiating objects from concrete subclasses. It is these objects that will be passed around your system.

Patterns for Organizing Objects and Classes

These patterns help you to organize the compositional relationships of your objects. More simply, these patterns show how you combine objects and classes.

Task-Oriented Patterns

These patterns describe the mechanisms by which classes and objects cooperate to achieve objectives.

Enterprise Patterns

I look at some patterns that describe typical Internet programming problems and solutions. Drawn largely from *Patterns of Enterprise Application Architecture* and *Core J2EE Patterns*, the patterns deal with presentation, and application logic.

Database Patterns

An examination of patterns that help with storing and retrieving data and with mapping objects to and from databases.

Summary

In this chapter, I examined some of the principles that underpin many design patterns. I looked at the use of composition to enable object combination and recombination at runtime, resulting in more flexible structures than would be available using inheritance alone. I introduced you to decoupling, the practice of extracting software components from their context to make them more generally applicable. I reviewed the importance of interface as a means of decoupling clients from the details of implementation.

In the coming chapters, I will examine some design patterns in detail.