

Notes on Software Patterns: Introduction to Patterns

H. Conrad Cunningham

16 April 2022

Contents

1	Introduction to Patterns	2
1.1	Chapter Introduction	2
1.2	What is a Pattern?	2
1.3	Descriptions of Patterns	3
1.3.1	Context	3
1.3.2	Problem	3
1.3.3	Solution	4
1.3.4	Aside: Extending pattern elements	4
1.4	Categories of Patterns	4
1.4.1	Architectural patterns	5
1.4.2	Design patterns	6
1.4.3	Idioms	8
1.5	What Next?	8
1.6	Acknowledgements	9
1.7	Terms and Concepts	9
1.8	References	9

Copyright (C) 1998-2022, H. Conrad Cunningham
Professor of Computer and Information Science
University of Mississippi
214 Weir
P.O. Box 1848
University, MS 38677
(662) 915-7396 (dept. office)

Browser Advisory: The HTML version of this textbook requires a browser that supports the display of MathML. A good choice as of April 2022 is a recent version of Firefox from Mozilla.

1 Introduction to Patterns

1.1 Chapter Introduction

The goal of this chapter is to introduce the basic concepts and terminology of software patterns as used in software architecture, software engineering, and programming.

The chapter approaches patterns mostly from the perspective of object-oriented programming using languages such as Java and Scala. Classic works on software “design patterns” include the Gamma et al. (i.e., the “Gang of Four”) [3] and Buschmann et al. (i.e., “Siemens”) [2] books. In these notes, we primarily use the terminology of Buschmann et al. [2].

Similar approaches may be used in functional languages such as Haskell, but often functional languages will use first-class and higher-order functions to express the patterns.

The accompanying set of HTML slides (not fully updated in 2022) is the following:

- Introduction to Patterns (HTML)

1.2 What is a Pattern?

When experts need to solve a problem, they seldom invent a totally new solution. More often they will recall a similar problem they have solved previously and reuse the essential aspects of the old solution to solve the new problem. They tend to think in problem-solution pairs.

Identifying the essential aspects of specific problem-solution pairs leads to descriptions of problem-solving *patterns* that can be reused.

The concept of a pattern as used in software architecture is borrowed from the field of (building) architecture, in particular from the writings of architect Christopher Alexander [1]. Buschmann et al. [2] defines pattern as follows in the context of software architecture:

“A *pattern for software architecture* describes a particular recurring design problem that arises in specific design contexts and presents a well-proven generic scheme for its solution. The solution scheme is specified by describing its constituent components, their responsibilities and relationships, and the ways in which they collaborate.”

Where software architecture is concerned, the concept of a pattern described here is essentially the same concept as an *architectural style* or *architectural idiom* in Shaw and Garlan [9].

In general, patterns have the following characteristics [2]:

- A pattern describes a solution to a recurring problem that arises in specific design situations.

- Patterns are not invented; they are distilled from practical experience.
- Patterns describe a group of components (e.g., classes or objects), how the components interact, and the responsibilities of each component. That is, they are higher level abstractions than classes or objects.
- Patterns provide a vocabulary for communication among designers. The choice of a name for a pattern is very important.
- Patterns help document the architectural vision of a design. If the vision is clearly understood, it will less likely be violated when the system is modified.
- Patterns provide a conceptual skeleton for a solution to a design problem and, hence, encourage the construction of software with well-defined properties.
- Patterns are building blocks for the construction of more complex designs.
- Patterns help designers manage the complexity of the software. When a recurring pattern is identified, the corresponding general solution can be implemented productively to provide a reliable software system.

1.3 Descriptions of Patterns

Various authors use different formats (i.e., “languages”) for describing patterns. Typically a pattern will be described with a schema that includes at least the following three elements [2]:

1. Context
2. Problem
3. Solution

1.3.1 Context

The *Context* element describes the circumstances in which the problem arises.

1.3.2 Problem

The *Problem* element describes the specific problem that arises repeatedly in the context.

In particular, the description describes the set of *forces* repeatedly arising in the context. A force is some aspect of the problem that must be considered when attempting a solution. Example types of forces include:

- requirements the solution must satisfy (e.g., efficiency)
- constraints that must be considered (e.g., use of a certain algorithm or protocol)

- desirable properties of a solution (e.g., easy to modify)

Forces may be complementary (i.e., can be achieved simultaneously) or contradictory (i.e., can only be balanced).

1.3.3 Solution

The *Solution* section describes a proven solution to the problem.

The solution specifies a configuration of elements to balance the forces associated with the problem.

- A pattern describes the static structure of the configuration, identifying the components and the connectors (i.e., the relationships among the components).
- A pattern also describes the dynamic runtime behavior of the configuration, identifying the control structure of the components and connectors.

1.3.4 Aside: Extending pattern elements

In a helpful tutorial on pattern writing, Wellhausen and Fiesser [13] separate out the Forces and Consequence as separate elements and state that the following elements must be present in the given order:

- *Pattern Name* gives an evocative name for the pattern.
- *Context* describes the circumstances in which the problem occurs.
- *Problem* describes the specific problem to be solved.
- *Forces* describe why the problem is difficult to solve, giving different considerations that must be balanced to solve the problem.
- *Solution* describes how the solution to the problem works at an appropriate level of detail.
- *Consequences* describe what happens when a software designer applies the pattern. It gives both the possible *benefits* and possible *liabilities* of using the pattern.

All of the above elements except Consequences are also prescribed by the MANDATORY ELEMENTS PRESENT pattern from Meszaros and Doble's *A Pattern Language for Pattern Writing* [4]. Their OPTIONAL ELEMENTS WHEN HELPFUL pattern suggests use of optional elements such as Examples, Code Samples, and Rationale. Following their READABLE REFERENCES TO PATTERNS pattern, we refer to a pattern using its name in small capital letters.

1.4 Categories of Patterns

Patterns can be grouped into three categories according to their level of abstraction [2]:

1. Architectural patterns
2. Design patterns
3. Idioms

1.4.1 Architectural patterns

Buschmann et al. [2] defines an architectural pattern as follows:

“An *architectural pattern* expresses a fundamental structural organization schema for software systems. It provides a set of predefined subsystems, specifies their responsibilities, and includes rules and guidelines for organizing the relationships between them.”

In early work, Shaw and Garlan [9] used the term *architectural style* instead of architectural pattern.

An architectural pattern is a high-level abstraction. The choice of the architectural pattern to be used is a fundamental design decision in the development of a software system. It determines the system-wide structure and constrains the design choices available for the various subsystems. It is, in general, independent of the implementation language used.

Examples of architectural patterns include the following.

- The PIPES AND FILTERS (or PIPELINE) pattern [2:53,8,9,15] defines a structure for systems in which an independent set of computations—called *filters*—transform one or more input streams—passing along *pipes*—incrementally to create one or more output streams.

See the separate notes on the Pipes and Filters pattern or the set of Powerpoint slides for more a more detailed discussion of this pattern.

In the UNIX operating system [6], for instance, a filter is a program that reads a stream of bytes from its standard input and writes a transformed stream to its standard output. These programs can be connected together (e.g., in the interactive shell program) with the output of one filter becoming the input of the next filter in the sequence via the pipe mechanism. Larger systems can thus be constructed from simple components that otherwise operate independently of one another.

- The LAYERED SYSTEMS [2:31,8,9] pattern organizes a system hierarchically with “each layer providing service to the layer above it and serving as a client for the layer below” [9].

Communication protocols, operating systems, virtual machines, and application programming interfaces (APIs) are often designed and implemented as layered systems [2].

- The BLACKBOARD pattern [2:71,8,9,16] defines a structure in which “a collection of independent programs”—the *knowledge sources*—“work cooperatively on a common data structure”—the *blackboard* [2].

A BLACKBOARD pattern is one of two subcategories of the REPOSITORY pattern [8,9]. In a BLACKBOARD system, the state of the blackboard (i.e., the repository) triggers the activity of the knowledge sources (i.e., the independent programs). Such a structure is often useful in artificial intelligence applications.

The other subcategory of the Repository pattern is for situations in which the types of the inputs trigger the independent programs. In such systems, the repository can be a traditional database that simply organizes the data collected.

- The MODEL-VIEW-CONTROLLER [2:125,14] pattern defines a structure for interactive user interface programs (such as GUIs or Web applications). It decomposes the application into three types of components [2]:
 - The Model contains the application’s core functionality (i.e., its logic) and data (e.g., accesses the application’s database).
 - A View represents the data contained in the Model for display to a user.
 - A Controller handles user inputs associated with a View and interacts with the View and Model to carry out the user’s commands.

An application has a single Model and one or more Views. Each View has a unique Controller to enable the user to manipulate both that View and the Model.

The Buschmann et al. book [2], Qian et al. book [5], and Shaw paper [9] elaborate on these and other architectural patterns. (**This link** is to a local copy of a preprint of the Shaw paper.)

1.4.2 Design patterns

Buschmann et al. [2] defines design pattern as follows:

“A *design pattern* provides a scheme for refining the subsystems or components of a software system, or the relationships between them. It describes a commonly-recurring structure of communicating components that solves a general design problem within a particular context.”

A design pattern is a mid-level abstraction. The choice of a design pattern does not affect the fundamental structure of the software system, but it does affect the structure of a subsystem. Like the architectural pattern, the design pattern tends to be independent of the implementation language to be used.

A design pattern might not, however, be independent of the programming

paradigm. A pattern that is meaningful for a statically typed, object-oriented language without first-class functions may not be as meaningful for a dynamically typed language or for a functional language with first-class functions.

Examples of design patterns include the following.

- The ADAPTER (or WRAPPER) pattern [3:139,10,17] converts the interface of one existing type of object to have the same interface as a different existing type of object.

For example, suppose a Java program has a base class `Stack` to represent stack data structures. Also suppose the program has instances of the builtin class `Vector` that we wish to use as `Stack` objects. We can implement a new subclass of `Stack`, say `VectorAsStack`, that wraps a `Vector` object and implements the operations of `Stack` by delegating them appropriately to the `Vector` object but hides the non-stack features of the `Vector`.

As another example, consider Schmid's paper "Creating Applications from Components: A Manufacturing Framework Design" [7]. The application framework presented in the paper uses the ADAPTER pattern. It adapts the portal robot machine class so that its instances can be used in place of transport service class instances.

- The ITERATOR pattern [3:257,11,18] defines a way to access the elements of a container (data structure) sequentially without exposing the container's representation.

Iterator objects for standard collections are now common in most programming language libraries. Programmers can also implement iterators for their own custom collections.

- The STRATEGY (or POLICY) pattern [3:315,12,19] defines an interface to a family of related algorithms so that any algorithm in the family can be dynamically substituted for another at runtime.

Suppose we have an container class `C` whose elements we wish to be able to sort using one of several different sorting algorithms selected at runtime. To apply the STRATEGY pattern, we design `C` so that it delegates the sorting of its elements to a method `sort()` on an object of type `Sorter` stored its instance variable `mySorter`. We then implement a different subclass of `Sorter` for each sorting algorithm of interest. We can dynamically change `C`'s sorting behavior by assigning a different `Sorter` object to `mySorter`.

We also see this pattern used in the Schmid article [7]. Schmid's third transformation involves breaking up the application logic class `ProcessingControl` into several subclasses of a new `ProcessingStrategy` class. The specific processing strategy can then be selected dynamically based on the specific part-processing task.

See the Powerpoint slides on the STRATEGY pattern for more information.

1.4.3 Idioms

Buschann et al. [2] defines idiom as follows:

“An *idiom* is a low-level pattern specific to a programming language. An idiom describes how to implement particular aspects of components or the relationships between them using the features of the given language.”

An idiom is a low-level abstraction. It is usually a language-specific pattern that deals with some aspects of both design and implementation.

In some sense, use of a consistent program coding and formatting style can be considered an idiom for the language being used. Such a style would provide guidelines for naming variables, laying out declarations, indenting control structures, ordering the features of a class, determining how values are returned, and so forth. A good style that is used consistently makes a program easier to understand than otherwise would be the case.

In Java, the language-specific iterator defined to implement the `Iterator` interface can be considered an idiom. It is a language-specific instance of the more general `ITERATOR` design pattern.

Another example of an idiom is the use of the `COUNTED POINTER` (or `COUNTED BODY` or `REFERENCE COUNTING`) technique for storage management of shared objects in C++. In this idiom, we control access to a shared object through two classes, a *Body* (representation) class and a *Handle* (access) class.

An object of the *Body* class holds the shared object and a count of the number of references to the object.

An object of a *Handle* class holds a direct reference to a body object; all other parts of the program must access the body indirectly through handle class methods. The handle methods can increment the reference count when a new reference is created and decrement the count when a reference is freed. When a reference count goes to zero, the shared object and its body can be deleted. Often the programmer using this pattern will want to override the `operator->` of the handle class to give more transparent access to the shared object.

We can use a variant of the `COUNTED POINTER` idiom to implement a “copy on write” mechanism. That is, the body is shared as long as only “read” access is needed, but a copy is created whenever one of the holders makes a change to the state of the object.

1.5 What Next?

TODO

1.6 Acknowledgements

I wrote the first version of these notes for my Spring 1998 graduate Special Topics in Software Architecture class. I based the notes, in part, on:

- Chapter 1 of the “Siemens” book [2]
- Chapter 2 of Shaw and Garlan [9]

I revised the notes somewhat for related courses in 1999, 2000, 2001, 2002, and 2004. Also, in 2004 I revised the notes, created slides, and included them as a part of the materials supported by a grant from the Acxiom Corporation titled “The Acxiom Laboratory for Software Architecture and Component Engineering (ALSACE)”. My ALSACE research team included PhD students Yi Liu and Pallavi Tadepalli and MS students Mingxian Fu and “Melody” Hui Xiong.

In Spring 2017, I adapted the earlier notes to use Pandoc-flavored Markdown. In Spring 2018 I revised the notes slightly to fit in with the other documents for the CSci 658 course.

I retired from the full-time faculty in May 2019. As one of my post-retirement projects, I am continuing work on possible textbooks based on the course materials I had developed during my three decades as a faculty member. In January 2022, I began refining the existing content, integrating separately developed materials together, reformatting the documents, constructing a unified bibliography (e.g., using citeproc), and improving my build workflow and use of Pandoc.

In 2022, I also added the aside based on Wellhausen and Fiesser’s tutorial and expanded the discussion of the example patterns.

I maintain this chapter as text in Pandoc’s dialect of Markdown using embedded LaTeX markup for the mathematical formulas and then translate the document to HTML, PDF, and other forms as needed

1.7 Terms and Concepts

TODO: Update

Pattern (for software architecture), architectural pattern (or architectural style), design pattern, idiom. Pattern Context, Problem, and Solution. Pattern Name, Forces, and Consequences. Pipes and Filters, Layers, Blackboard, and Model-View-Controller architectural patterns. Adapter, Iterator, and Strategy design patterns. Counted Pointer (or Reference Counting, Handle-Body) idiom.

1.8 References

- [1] Christopher Alexander. 1977. *A pattern language: Towns, buildings, construction*. Oxford University Press, Oxford, UK.

- [2] Frank Buschmann, Regine Meunier, Hans Rohnert, Peter Sommerlad, and Michael Stal. 1996. *Pattern-oriented software architecture: A system of patterns*. Wiley, Hoboken, New Jersey, USA.
- [3] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. 1995. *Design patterns: Elements of reusable object-oriented software*. Addison-Wesley, Boston, Massachusetts, USA.
- [4] Gerard Meszaros and Jim Doble. 1998. A pattern language for pattern writing. In *Pattern languages of program design 3*, Addison-Wesley, Boston, Massachusetts, USA, 529–574.
- [5] Kai Qian, Xiang Fu, Lixin Tao, Chong-Wei Xu, and Jorge L. Diaz-Herrera. 2010. *Software architecture and design illuminated*. Jones & Bartlett Learning, Burlington, Massachusetts, USA.
- [6] Dennis M. Ritchie. 1984. The UNIX system: The evolution of the UNIX time-sharing system. *AT&T Bell Laboratories Technical Journal* 63, 8 (1984), 1577–1593.
- [7] Hans Albrecht Schmid. 1996. Creating applications from components: A manufacturing framework design. *IEEE Software* 13, 6 (1996), 67–75.
- [8] Mary Shaw. 1996. Some patterns for software architecture. In *Pattern languages of program design 2*, John M. Vlissides, James O. Coplien and Norman L. Kerth (eds.). Addison-Wesley, Boston, Massachusetts, USA, 255–270.
- [9] Mary Shaw and David Garlan. 1996. *Software architecture: Perspectives on an emerging discipline*. Prentice-Hall, Englewood Cliffs, New Jersey, USA.
- [10] Source Making. 2022. Adapter design pattern. Retrieved from https://sourcemaking.com/design_patterns/adapter
- [11] Source Making. 2022. Iterator design pattern. Retrieved from https://sourcemaking.com/design_patterns/iterator
- [12] Source Making. 2022. Strategy design pattern. Retrieved from https://sourcemaking.com/design_patterns/strategy
- [13] Tim Wellhausen and Andreas Fiesser. 2011. How to write a pattern? A rough guide for first-time pattern authors. In *Proceedings of the 16th Europe an conference on pattern languages of programs (EuroPLOP '11)*, Irsee, Germany, 1–9 (Article 5).
- [14] Wikipedia: The Free Encyclopedia. 2022. Model-view-controller. Retrieved from <https://en.wikipedia.org/wiki/Model-view-controller>
- [15] Wikipedia: The Free Encyclopedia. 2022. Pipeline (software). Retrieved from [https://en.wikipedia.org/wiki/Pipeline_\(software\)](https://en.wikipedia.org/wiki/Pipeline_(software))
- [16] Wikipedia: The Free Encyclopedia. 2022. Blackboard (design pattern). Retrieved from [https://en.wikipedia.org/wiki/Blackboard_\(design_pattern\)](https://en.wikipedia.org/wiki/Blackboard_(design_pattern))

- [17] Wikipedia: The Free Encyclopedia. 2022. Adapter pattern. Retrieved from https://en.wikipedia.org/wiki/Adapter_pattern
- [18] Wikipedia: The Free Encyclopedia. 2022. Iterator pattern. Retrieved from https://en.wikipedia.org/wiki/Iterator_pattern
- [19] Wikipedia: The Free Encyclopedia. 2022. Strategy pattern. Retrieved from https://en.wikipedia.org/wiki/Strategy_pattern