



PRACTICAL OBJECT-ORIENTED DESIGN

AN AGILE PRIMER USING RUBY

SECOND EDITION



SANDI METZ

Praise for the first edition of *Practical Object-Oriented Design in Ruby*

“Meticulously pragmatic and exquisitely articulate, Practical Object Oriented Design in Ruby makes otherwise elusive knowledge available to an audience which desperately needs it. The prescriptions are appropriate both as rules for novices and as guidelines for experienced professionals.”

—Katrina Owen, Creator, Exercism

“I do believe this will be the most important Ruby book of 2012. Not only is the book 100% on-point, Sandi has an easy writing style with lots of great analogies that drive every point home.”

—Avdi Grimm, author of *Exceptional Ruby and Objects on Rails*

“While Ruby is an object-oriented language, little time is spent in the documentation on what OO truly means or how it should direct the way we build programs. Here Metz brings it to the fore, covering most of the key principles of OO development and design in an engaging, easy-to-understand manner. This is a must for any respectable Ruby bookshelf.”

—Peter Cooper, editor, *Ruby Weekly*

“So good, I couldn’t put it down! This is a must-read for anyone wanting to do object-oriented programming in any language, not to mention it has completely changed the way I approach testing.”

—Charles Max Wood, Ruby Rogues Podcast co-host and CEO of Devchat.tv

“Distilling scary OO design practices with clear-cut examples and explanations makes this a book for novices and experts alike. It is well worth the study by anyone interested in OO design being done right and ‘light.’ I thoroughly enjoyed this book.”

—Manuel Pais, DevOps and Continuous Delivery Consultant, Independent

“If you call yourself a Ruby programmer, you should read this book. It’s jam-packed with great nuggets of practical advice and coding techniques that you can start applying immediately in your projects.”

—Ylan Segal, San Diego Ruby User Group

“This is the best OO book I’ve ever read. It’s short, sweet, but potent. It slowly moves from simple techniques to more advanced, each example improving on the last. The ideas it presents are useful not just in Ruby but in static languages like C# too. Highly recommended!”

—Kevin Berridge, software engineering manager,
Pointe Blank Solutions, and organizer, Burning River Developers Meetup

“This is the best programming book I’ve read in ages. Sandi talks about basic principles, but these are things we’re probably still doing wrong and she shows us why and how. The book has the perfect mix of code, diagrams, and words. I can’t recommend it enough and if you’re serious about being a better programmer, you’ll read it and agree.

—Derick Hitchcock, software engineer, Cisco

“Metz’s take on the subject is rooted strongly in theory, but the explanation always stays grounded in real world concerns, which helped me to internalize it. The book is clear and concise, yet achieves a tone that is more friendly than terse.”

—Alex Strasheim, network administrator, Ensemble Travel Group

“Whether you’re just getting started in your software development career, or you’ve been coding for years (like I have), it’s likely that you’ll learn a lot from Ms. Metz’s book. She does a fantastic job of explaining the whys of well-designed software along with the hows.”

—Gabe Hollombe, software craftsman, avantbard.com

PRACTICAL OBJECT-ORIENTED DESIGN

Second Edition

This page intentionally left blank

PRACTICAL OBJECT-ORIENTED DESIGN

An Agile Primer Using Ruby

Second Edition

Sandi Metz

◆ Addison-Wesley

Boston • Columbus • New York • San Francisco • Amsterdam • Cape Town • Dubai
London • Madrid • Milan • Munich • Paris • Montreal • Toronto • Delhi • Mexico City
São Paulo • Sydney • Hong Kong • Seoul • Singapore • Taipei • Tokyo

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Visit us on the Web: informit.com/aw

Library of Congress Control Number: 2018939833

Copyright © 2019 Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearsoned.com/permissions/.

ISBN-13: 978-0-13-445647-8

ISBN-10: 0-13-445647-5

ScoutAutomatedPrintCode

For Amy, who read everything first

This page intentionally left blank

Contents

Introduction	xv
Acknowledgments	xix
About the Author	xxi
1 Object-Oriented Design	1
1.1 In Praise of Design	2
1.1.1 The Problem Design Solves	2
1.1.2 Why Change Is Hard	3
1.1.3 A Practical Definition of Design	3
1.2 The Tools of Design	4
1.2.1 Design Principles	4
1.2.2 Design Patterns	6
1.3 The Act of Design	6
1.3.1 How Design Fails	6
1.3.2 When to Design	7
1.3.3 Judging Design	9
1.4 A Brief Introduction to Object-Oriented Programming	10
1.4.1 Procedural Languages	11
1.4.2 Object-Oriented Languages	11
1.5 Summary	13
2 Designing Classes with a Single Responsibility	15
2.1 Deciding What Belongs in a Class	16
2.1.1 Grouping Methods into Classes	16
2.1.2 Organizing Code to Allow for Easy Changes	16
2.2 Creating Classes That Have a Single Responsibility	17
2.2.1 An Example Application: Bicycles and Gears	17
2.2.2 Why Single Responsibility Matters	21

2.2.3	Determining If a Class Has a Single Responsibility	22
2.2.4	Determining When to Make Design Decisions	22
2.3	Writing Code That Embraces Change	24
2.3.1	Depend on Behavior, Not Data	24
2.3.2	Enforce Single Responsibility Everywhere	29
2.4	Finally, the Real Wheel	33
2.5	Summary	35
3	Managing Dependencies	37
3.1	Understanding Dependencies	38
3.1.1	Recognizing Dependencies	39
3.1.2	Coupling Between Objects (CBO)	39
3.1.3	Other Dependencies	40
3.2	Writing Loosely Coupled Code	41
3.2.1	Inject Dependencies	41
3.2.2	Isolate Dependencies	44
3.2.3	Remove Argument-Order Dependencies	48
3.3	Managing Dependency Direction	53
3.3.1	Reversing Dependencies	53
3.3.2	Choosing Dependency Direction	55
3.4	Summary	59
4	Creating Flexible Interfaces	61
4.1	Understanding Interfaces	61
4.2	Defining Interfaces	63
4.2.1	Public Interfaces	64
4.2.2	Private Interfaces	64
4.2.3	Responsibilities, Dependencies, and Interfaces	64
4.3	Finding the Public Interface	65
4.3.1	An Example Application: Bicycle Touring Company	65
4.3.2	Constructing an Intention	65
4.3.3	Using Sequence Diagrams	66
4.3.4	Asking for “What” Instead of Telling “How”	70
4.3.5	Seeking Context Independence	72
4.3.6	Trusting Other Objects	74
4.3.7	Using Messages to Discover Objects	75
4.3.8	Creating a Message-Based Application	77
4.4	Writing Code That Puts Its Best (Inter)Face Forward	77
4.4.1	Create Explicit Interfaces	77
4.4.2	Honor the Public Interfaces of Others	79

4.4.3	Exercise Caution When Depending on Private Interfaces	80
4.4.4	Minimize Context	80
4.5	The Law of Demeter	80
4.5.1	Defining Demeter	81
4.5.2	Consequences of Violations	81
4.5.3	Avoiding Violations	82
4.5.4	Listening to Demeter	83
4.6	Summary	84
5	Reducing Costs with Duck Typing	85
5.1	Understanding Duck Typing	85
5.1.1	Overlooking the Duck	86
5.1.2	Compounding the Problem	88
5.1.3	Finding the Duck	90
5.1.4	Consequences of Duck Typing	94
5.2	Writing Code That Relies on Ducks	95
5.2.1	Recognizing Hidden Ducks	95
5.2.2	Placing Trust in Your Ducks	97
5.2.3	Documenting Duck Types	98
5.2.4	Sharing Code between Ducks	98
5.2.5	Choosing Your Ducks Wisely	98
5.3	Conquering a Fear of Duck Typing	100
5.3.1	Subverting Duck Types with Static Typing	100
5.3.2	Static versus Dynamic Typing	101
5.3.3	Embracing Dynamic Typing	102
5.4	Summary	103
6	Acquiring Behavior through Inheritance	105
6.1	Understanding Classical Inheritance	105
6.2	Recognizing Where to Use Inheritance	106
6.2.1	Starting with a Concrete Class	107
6.2.2	Embedding Multiple Types	109
6.2.3	Finding the Embedded Types	111
6.2.4	Choosing Inheritance	112
6.2.5	Drawing Inheritance Relationships	114
6.3	Misapplying Inheritance	114
6.4	Finding the Abstraction	116
6.4.1	Creating an Abstract Superclass	117
6.4.2	Promoting Abstract Behavior	120
6.4.3	Separating Abstract from Concrete	123

6.4.4	Using the Template Method Pattern	125
6.4.5	Implementing Every Template Method	127
6.5	Managing Coupling between Superclasses and Subclasses	129
6.5.1	Understanding Coupling	129
6.5.2	Decoupling Subclasses Using Hook Messages	134
6.6	Summary	139
7	Sharing Role Behavior with Modules	141
7.1	Understanding Roles	142
7.1.1	Finding Roles	142
7.1.2	Organizing Responsibilities	143
7.1.3	Removing Unnecessary Dependencies	146
7.1.4	Writing the Concrete Code	147
7.1.5	Extracting the Abstraction	150
7.1.6	Looking Up Methods	153
7.1.7	Inheriting Role Behavior	157
7.2	Writing Inheritable Code	158
7.2.1	Recognize the Antipatterns	158
7.2.2	Insist on the Abstraction	158
7.2.3	Honor the Contract	159
7.2.4	Use the Template Method Pattern	160
7.2.5	Preemptively Decouple Classes	160
7.2.6	Create Shallow Hierarchies	160
7.3	Summary	161
8	Combining Objects with Composition	163
8.1	Composing a Bicycle of Parts	163
8.1.1	Updating the Bicycle Class	164
8.1.2	Creating a Parts Hierarchy	165
8.2	Composing the Parts Object	168
8.2.1	Creating a Part	168
8.2.2	Making the Parts Object More Like an Array	172
8.3	Manufacturing Parts	176
8.3.1	Creating the PartsFactory	177
8.3.2	Leveraging the PartsFactory	179
8.4	The Composed Bicycle	181
8.5	Deciding between Inheritance and Composition	185
8.5.1	Accepting the Consequences of Inheritance	186
8.5.2	Accepting the Consequences of Composition	188
8.5.3	Choosing Relationships	189
8.6	Summary	191

9 Designing Cost-Effective Tests 193

- 9.1 Intentional Testing 194
 - 9.1.1 Knowing Your Intentions 194
 - 9.1.2 Knowing What to Test 196
 - 9.1.3 Knowing When to Test 199
 - 9.1.4 Knowing How to Test 200
- 9.2 Testing Incoming Messages 202
 - 9.2.1 Deleting Unused Interfaces 204
 - 9.2.2 Proving the Public Interface 204
 - 9.2.3 Isolating the Object under Test 206
 - 9.2.4 Injecting Dependencies Using Classes 208
 - 9.2.5 Injecting Dependencies as Roles 210
- 9.3 Testing Private Methods 215
 - 9.3.1 Ignoring Private Methods during Tests 216
 - 9.3.2 Removing Private Methods from the Class under Test 216
 - 9.3.3 Choosing to Test a Private Method 216
- 9.4 Testing Outgoing Messages 217
 - 9.4.1 Ignoring Query Messages 217
 - 9.4.2 Proving Command Messages 218
- 9.5 Testing Duck Types 221
 - 9.5.1 Testing Roles 221
 - 9.5.2 Using Role Tests to Validate Doubles 227
- 9.6 Testing Inherited Code 233
 - 9.6.1 Specifying the Inherited Interface 233
 - 9.6.2 Specifying Subclass Responsibilities 236
 - 9.6.3 Testing Unique Behavior 240
- 9.7 Summary 244

Afterword 245**Index 247**

This page intentionally left blank

Introduction

We want to do our best work, and we want the work we do to have meaning. And, all else being equal, we prefer to enjoy ourselves along the way.

Those of us whose work is to write software are incredibly lucky. Building software is a guiltless pleasure because we get to use our creative energy to get things done. We have arranged our lives to have it both ways; we can enjoy the pure act of writing code in sure knowledge that the code we write has use. We produce things that matter. We are modern craftspeople, building structures that make up present-day reality, and no less than bricklayers or bridge builders, we take justifiable pride in our accomplishments.

This all programmers share, from the most enthusiastic newbie to the apparently jaded elder, whether working at the lightest weight Internet startup or the most staid, long-entrenched enterprise. We want to do our best work. We want our work to have meaning. We want to have fun along the way.

And so it's especially troubling when software goes awry. Bad software impedes our purpose and interferes with our happiness. Where once we felt productive, now we feel thwarted. Where once fast, now slow. Where once peaceful, now frustrated.

This frustration occurs when it costs too much to get things done. Our internal calculators are always running, comparing total amount accomplished to overall effort expended. When the cost of doing work exceeds its value, our efforts feel wasted. If programming gives joy it is because it allows us to be useful, when it becomes painful it is a sign that we believe we could, and should, be doing more. Our pleasure follows in the footsteps of work.

This book is about designing object-oriented software. It is not an academic tome, it is a programmer's story about how to write code. It teaches how to arrange software so as to be productive today and to remain so next month and next year. It shows how to write applications that can succeed in the present and still adapt to the future. It allows you to raise your productivity and reduce your costs for the entire lifetime of your applications.

This book believes in your desire to do good work and gives you the tools you need to best be of use. It is completely practical and as such is, at its core, a book about how to write code that brings you joy.

Who Might Find This Book Useful?

This book assumes that you have at least tried to write object-oriented software. It is not necessary that you feel you succeeded, just that you made the attempt in any object-oriented (OO) language. Chapter 1, “Object-Oriented Design,” contains a brief overview of object-oriented programming (OOP), but its goal is to define common terms, not to teach programming.

If you want to learn OO design (OOD) but have not yet done any object-oriented programming, at least take a tutorial before reading this book. OOD solves problems; suffering from those problems is very nearly a prerequisite for comprehending these solutions. Experienced programmers may be able to skip this step, but most readers will be happier if they write some OO code before starting this book.

This book uses Ruby to teach OOD but you do not need to know Ruby to understand the concepts herein. There are many code examples but all are quite straightforward. If you have programmed in any OO language you will find Ruby easy to understand.

If you come from a statically typed OO language like Java or C++ you have the background necessary to benefit from reading this book. The fact that Ruby is dynamically typed simplifies the syntax of the examples and distills the design ideas to their essence, but every concept in this book can be directly translated to a statically typed OO language.

How to Read This Book

Chapter 1 contains a general overview of the whys, whens, and wherefores of OO design, followed by a brief overview of object-oriented programming. This chapter stands alone. You can read it first, last, or, frankly, skip it entirely, although if you are currently stuck with an application that suffers from lack of design, you may find it a comforting tale.

If you have experience writing object-oriented applications and want to jump right in, you can safely start with Chapter 2. If you do so and then stumble upon an unfamiliar term, come back and browse the “Introduction to Object-Oriented Programming” section of Chapter 1, which introduces and defines common OO terms used throughout the book.

Chapters 2 through 9 progressively explain object-oriented design. Chapter 2, “Designing Classes with a Single Responsibility,” covers how to decide what belongs in a single class. Chapter 3, “Managing Dependencies,” illustrates how objects get entangled with one another and shows how to keep them apart. These two chapters are focused on objects rather than messages.

In Chapter 4, “Creating Flexible Interfaces,” the emphasis begins to shift away from object-centric toward message-centric design. Chapter 4 is about defining interfaces and is concerned with how objects talk to one another. Chapter 5, “Reducing Costs with Duck Typing,” is about duck typing and introduces the idea that objects of different classes may play common roles. Chapter 6, “Acquiring Behavior through Inheritance,” teaches the techniques of classical inheritance, which are then used in Chapter 7, “Sharing Role Behavior with Modules,” to create duck typed roles. Chapter 8, “Combining Objects with Composition,” explains the technique of building objects via composition and provides guidelines for choosing among composition, inheritance, and duck-typed role sharing. Chapter 9, “Designing Cost-Effective Tests,” concentrates on the design of tests, which it illustrates using code from earlier chapters of the book.

Each of these chapters builds on the concepts of the last. They are full of code and best read in order.

How to Use This Book

This book will mean different things to readers of different backgrounds. Those already familiar with OOD will find things to think about, possibly encounter some new points of view, and probably disagree with a few of the suggestions. Because there is no final authority on OOD, challenges to the principles (and to this author) will improve the understanding of all. In the end, you must be the arbiter of your own designs; it is up to you to question, to experiment, and to choose.

While this book should be of interest to many levels of reader, it is written with the particular goal of being accessible to novices. If you are one of those novices, this part of the introduction is especially for you. Know this: Object-oriented design is not black magic. It is simply things you don't yet know. The fact that you've read this far indicates you care about design; this desire to learn is the only prerequisite for benefiting from this book.

Chapters 2 through 9 explain OOD principles and provide very explicit programming rules; these rules will mean different things to novices than they mean to experts. If you are a novice, start out by following these rules in blind faith if necessary. This early obedience will stave off disaster until you can gain enough experience to make

your own decisions. By the time the rules start to chafe, you'll have enough experience to make up rules of your own, and your career as a designer will have begun.

Software Versions Used in This Book

The examples in this book were written using Ruby 2.4 and tested with Minitest 5.10.3. Source code for the examples can be found at <https://github.com/skmetz/poodr2>.

Register your copy of *Practical Object-Oriented Design, Second Edition*, on the InformIT site for convenient access to updates and/or corrections as they become available. To start the registration process, go to informit.com/register and log in or create an account. Enter the product ISBN (9780134456478) and click Submit. Look on the Registered Products tab for an Access Bonus Content link next to this product, and follow that link to access any available bonus materials. If you would like to be notified of exclusive offers on new editions and updates, please check the box to receive email from us.

Acknowledgments

It is a wonder this book exists; the fact that it does is due to the efforts and encouragement of many people.

Throughout the long process of writing, Lori Evans and TJ Stankus provided early feedback on every chapter. They live in Durham, NC, and thus could not escape me, but this fact does nothing to lessen my appreciation for their help.

Midway through the book, after it became impossible to deny that its writing would take approximately twice as long as originally estimated, Mike Dalessio and Gregory Brown read drafts and gave invaluable feedback and support. Their encouragement and enthusiasm kept the project alive during dark days.

A number of reviewers cast their keen eyes on the entire book, acting as gracious stand-ins for you, the gentle reader. As the first edition neared completion, Steve Klabnik, Desi McAdam, and Seth Wax gave it careful readings. The second edition was meticulously scoured by Will Sommers and Tory Peterschild. Their impressions and suggestions caused changes that will benefit all who follow.

Late drafts were given careful, thorough readings by Katrina Owen, Avdi Grimm, and Rebecca Wirfs-Brock, and the book is much improved by their kind and thoughtful feedback. Before they pitched in, Katrina, Avdi, and Rebecca were strangers to me; I am grateful for their involvement and humbled by their generosity. If you find this book useful, thank them when you next see them.

I am also grateful for the Gotham Ruby Group and for everyone who expressed their appreciation for the design talks I gave at GoRuCo 2009 and 2011. The folks at GoRuCo took a chance on an unknown and gave me a forum in which to express these ideas; this book started there. Ian McFarland and Brian Ford watched those talks and their immediate and ongoing enthusiasm for this project was both infectious and convincing.

The process of writing was greatly aided by Michael Thurston of Pearson, who was like an ocean liner of calmness and organization chugging through the chaotic sea of my opposing rogue writing waves. You can, I expect, see the problem he faced. He insisted, with endless patience and grace, that the writing be arranged in a readable structure. I believe his efforts have paid off and hope you will agree.

My thanks also to Debra Williams Cauley, my editor at Pearson/Addison-Wesley, who overheard an ill-timed hallway rant in 2006 at the first Ruby on Rails conference in Chicago and launched the campaign that eventually resulted in this book. Despite my best efforts, she would not take no for an answer. She cleverly moved from one argument to the next until she finally found the one that convinced; this accurately reflects her persistence and dedication.

I owe a debt to the entire object-oriented design community. I did not make up the ideas in this book, I am merely a translator, and I stand on the shoulders of giants. It goes without saying that while all credit for these ideas belongs to others—failures of translation are mine alone.

And finally, this book owes its existence to my partner Amy Germuth. Before this project started I could not imagine writing a book; her view of the world as a place where people did such things made doing so seem possible. The book in your hands is a tribute to her boundless patience and endless support.

Thank you, each and every one.

About the Author

Sandi Metz, author of *Practical Object-Oriented Design in Ruby* and *99 Bottles of OOP*, believes in simple code and straightforward explanations. She prefers working software, practical solutions, and lengthy bicycle trips (not necessarily in that order), and writes, consults, speaks, and teaches about object-oriented design.

This page intentionally left blank

CHAPTER 3

Managing Dependencies

Object-oriented programming languages contend that they are efficient and effective because of the way they model reality. Objects reflect qualities of a real-world problem and the interactions between those objects provide solutions. These interactions are inescapable. A single object cannot know everything, so inevitably it will have to talk to another object.

If you could peer into a busy application and watch the messages as they pass, the traffic might seem overwhelming. There's a lot going on. However, if you stand back and take a global view, a pattern becomes obvious. Each message is initiated by an object to invoke some bit of behavior. All of the behavior is dispersed among the objects. Therefore, for any desired behavior, an object either knows it personally, inherits it, or knows another object who knows it.

The previous chapter concerned itself with the first of these, that is, behaviors that a class should personally implement. The second, inheriting behavior, will be covered in Chapter 6, "Acquiring Behavior through Inheritance." This chapter is about the third, getting access to behavior when that behavior is implemented in *other* objects.

Because well-designed objects have a single responsibility, their very nature requires that they collaborate to accomplish complex tasks. This collaboration is powerful and perilous. To collaborate, an object must know something about others. *Knowing* creates a dependency. If not managed carefully, these dependencies will strangle your application.

3.1 Understanding Dependencies

An object depends on another object if, when one object changes, the other might be forced to change in turn.

Here's a modified version of the Gear class, where Gear is initialized with four familiar arguments. The gear_inches method uses two of them, rim and tire, to create a new instance of Wheel. Wheel has not changed since you last saw it in Chapter 2, "Designing Classes with a Single Responsibility."

Listing 3.1

```
1 class Gear
2   attr_reader :chainring, :cog, :rim, :tire
3   def initialize(chainring, cog, rim, tire)
4     @chainring = chainring
5     @cog       = cog
6     @rim       = rim
7     @tire      = tire
8   end
9
10  def gear_inches
11    ratio * Wheel.new(rim, tire).diameter
12  end
13
14  def ratio
15    chainring / cog.to_f
16  end
17  # ...
18 end
19
20 class Wheel
21   attr_reader :rim, :tire
22   def initialize(rim, tire)
23     @rim = rim
24     @tire = tire
25   end
26
27   def diameter
28     rim + (tire * 2)
29   end
30  # ...
31 end
32
33 puts Gear.new(52, 11, 26, 1.5).gear_inches
34 # => 137.0909090909091
```

Examine the preceding code and make a list of the situations in which Gear would be forced to change because of a change to Wheel. This code seems innocent, but it's sneakily complex. Gear has at least four dependencies on Wheel, enumerated as follows. Most of the dependencies are unnecessary; they are a side effect of the coding style. Gear does not need them to do its job. Their very existence weakens Gear and makes it harder to change.

3.1.1 Recognizing Dependencies

An object has a dependency when it knows:

- The name of another class. Gear expects a class named Wheel to exist.
- The name of a message that it intends to send to someone other than `self`. Gear expects a Wheel instance to respond to `diameter`.
- The arguments that a message requires. Gear knows that `Wheel.new` requires a `rim` and a `tire`.
- The order of those arguments. Gear knows that Wheel takes *positional* arguments and that the first should be `rim`, the second, `tire`.

Each of these dependencies creates a chance that Gear will be forced to change because of a change to Wheel. Some degree of dependency between these two classes is inevitable; after all, they *must* collaborate, but most of the dependencies listed above are unnecessary. These unnecessary dependencies make the code less *reasonable*. Because they increase the chance that Gear will be forced to change, these dependencies turn minor code tweaks into major undertakings where small changes cascade through the application, forcing many changes.

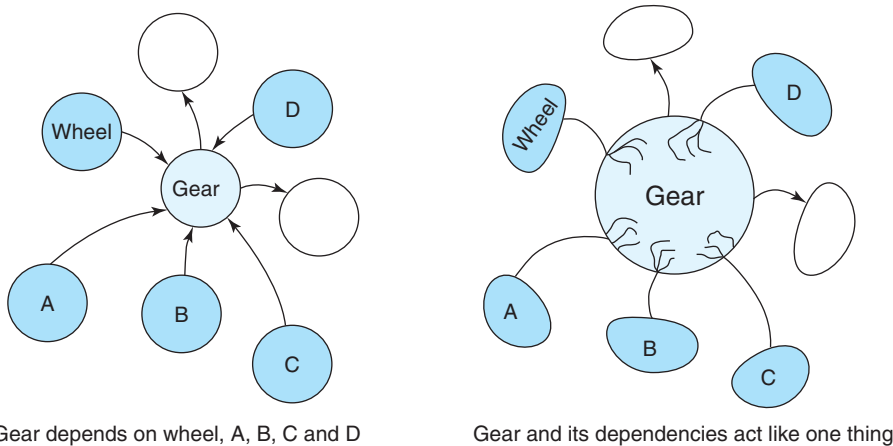
Your design challenge is to manage dependencies so that each class has the fewest possible; a class should know just enough to do its job and not one thing more.

3.1.2 Coupling Between Objects (CBO)

These dependencies *couple* Gear to Wheel. Alternatively, you could say that each coupling *creates* a dependency. The more Gear knows about Wheel, the more tightly coupled they are. The more tightly coupled two objects are, the more they behave like a single entity.

If you make a change to Wheel, you may find it necessary to make a change to Gear. If you want to reuse Gear, Wheel comes along for the ride. When you test Gear, you'll be testing Wheel too.

Figure 3.1 illustrates the problem. In this case, Gear depends on Wheel and four other objects, coupling Gear to five different things. When the underlying code was



Gear depends on wheel, A, B, C and D

Gear and its dependencies act like one thing

Figure 3.1 Dependencies entangle objects with one another

first written, everything worked fine. The problem lies dormant until you attempt to use `Gear` in another context or to change one of the classes upon which `Gear` depends. When that day comes, the cold hard truth is revealed; despite appearances, `Gear` is not an independent entity. Each of its dependencies is a place where another object is stuck to it. The dependencies cause these objects to act like a single thing. They move in lockstep; they change together.

When two (or three or more) objects are so tightly coupled that they behave as a unit, it's impossible to reuse just one. Changes to one object force changes to all. Left unchecked, unmanaged dependencies cause an entire application to become an entangled mess. A day will come when it's easier to rewrite everything than to change anything.

3.1.3 Other Dependencies

The remainder of this chapter examines the four kinds of dependencies listed previously and suggests techniques for avoiding the problems they create. However, before going forward, it's worth mentioning a few other common dependency-related issues that will be covered in other chapters.

One especially destructive kind of dependency occurs where an object knows another who knows another who knows something; that is, where many messages are chained together to reach behavior that lives in a distant object. This is the “knowing the name of a message you plan to send to someone other than *self*” dependency, only magnified. Message chaining creates a dependency between the original object and every object and message along the way to its ultimate target. These additional couplings greatly increase the chance that the first object will be forced to change because a change to *any* of the intermediate objects might affect it.

This case, a Law of Demeter violation, gets its own special treatment in Chapter 4, “Creating Flexible Interfaces.”

Another entire class of dependencies is that of tests on code. In the world outside of this book, tests come first. They drive design. However, they refer to code and thus depend on code. The natural tendency of “new-to-testing” programmers is to write tests that are too tightly coupled to code. This tight coupling leads to incredible frustration; the tests break every time the code is refactored, even when the fundamental behavior of the code does not change. Tests begin to seem costly relative to their value. Test-to-code over-coupling has the same consequence as code-to-code over-coupling. These couplings are dependencies that cause changes to the code to cascade into the tests, forcing them to change in turn.

The design of tests is examined in Chapter 9, “Designing Cost-Effective Tests.”

Despite these cautionary words, your application is not doomed to drown in unnecessary dependencies. As long as you recognize them, avoidance is quite simple. The first step to this brighter future is to understand dependencies in more detail; therefore, it’s time to look at some code.

3.2 Writing Loosely Coupled Code

Every dependency is like a little dot of glue that causes your class to stick to the things it touches. A few dots are necessary, but apply too much glue, and your application will harden into a solid block. Reducing dependencies means recognizing and removing the ones you don’t need.

The following examples illustrate coding techniques that reduce dependencies by decoupling code.

3.2.1 Inject Dependencies

Referring to another class by its name creates a major sticky spot. In the version of Gear we’ve been discussing (repeated in Listing 3.2), the `gear_inches` method contains an explicit reference to class `Wheel`.

Listing 3.2

```
1 class Gear
2   attr_reader :chainring, :cog, :rim, :tire
3   def initialize(chainring, cog, rim, tire)
4     @chainring = chainring
5     @cog       = cog
6     @rim       = rim
7     @tire      = tire
8   end
```

```
9
10 def gear_inches
11     ratio * Wheel.new(rim, tire).diameter
12 end
13 # ...
14 end
15
16 puts Gear.new(52, 11, 26, 1.5).gear_inches
17 # => 137.0909090909091
```

The immediate, obvious consequence of this reference is that if the name of the `Wheel` class changes, `Gear`'s `gear_inches` method must also change.

On the face of it, this dependency seems innocuous. After all, if a `Gear` needs to talk to a `Wheel`, something, somewhere, must create a new instance of the `Wheel` class. If `Gear` itself knows the name of the `Wheel` class, the code in `Gear` must be altered if `Wheel`'s name changes.

In truth, dealing with the name change is a relatively minor issue. You likely have a tool that allows you to do a global find/replace within a project. If `Wheel`'s name changes to `Wheely`, finding and fixing all of the references isn't that hard. However, the fact that line 11 above must change if the name of the `Wheel` class changes is the least of the problems with this code. A deeper problem exists that is far less visible but significantly more destructive.

When `Gear` hard-codes a reference to `Wheel` deep inside its `gear_inches` method, it is explicitly declaring that it is only willing to calculate gear inches for instances of `Wheel`. `Gear` refuses to collaborate with any other kind of object, even if that object has a diameter and uses gears.

If your application expands to include objects such as disks or cylinders and you need to know the gear inches of gears which use them, *you cannot*. Despite the fact that disks and cylinders naturally have a diameter, you can never calculate their gear inches because `Gear` is stuck to `Wheel`.

The code above exposes an unjustified attachment to type. It is not the class of the object that's important, it's the *message* you plan to send to it. `Gear` needs access to an object that can respond to `diameter`; a duck type, if you will (see Chapter 5, "Reducing Costs with Duck Typing"). `Gear` does not care and should not know about the class of that object. It is not necessary for `Gear` to know about the existence of the `Wheel` class in order to calculate `gear_inches`. It doesn't need to know that `Wheel` expects to be initialized with a `rim` and then a `tire`; it just needs an object that knows `diameter`.

Hanging these unnecessary dependencies on `Gear` simultaneously reduces `Gear`'s reusability and increases its susceptibility to being forced to change unnecessarily. `Gear` becomes less useful when it knows too much about *other* objects; if it knew less, it could do more.

Instead of being glued to `Wheel`, this next version of `Gear` expects to be initialized with an object that can respond to `diameter`:

Listing 3.3

```
1 class Gear
2   attr_reader :chainring, :cog, :wheel
3   def initialize(chainring, cog, wheel)
4     @chainring = chainring
5     @cog       = cog
6     @wheel     = wheel
7   end
8
9   def gear_inches
10    ratio * wheel.diameter
11  end
12  # ...
13 end
14
15 # Gear expects a 'Duck' that knows 'diameter'
16 puts Gear.new(52, 11, Wheel.new(26, 1.5)).gear_inches
17 # => 137.0909090909091
```

`Gear` now uses the `@wheel` variable to hold, and the `wheel` method to access, this object, but don't be fooled: `Gear` doesn't know or care that the object might be an instance of class `Wheel`. `Gear` only knows that it holds an object that responds to `diameter`.

This change is so small it is almost invisible, but coding in this style has huge benefits. Moving the creation of the new `Wheel` instance outside of `Gear` decouples the two classes. `Gear` can now collaborate with any object that implements `diameter`. As an extra bonus, this benefit was free. Not only is the resulting `Gear` class smaller than the original, but the decoupling was achieved by simply rearranging existing code.

This technique is known as *dependency injection*. Despite its daunting reputation, dependency injection truly is this simple. `Gear` previously had explicit dependencies on the `Wheel` class and on the type and order of its initialization arguments, but through injection these dependencies have been reduced to a single dependency on the `diameter` method. `Gear` is now smarter because it knows less.

Using dependency injection to shape code relies on your ability to recognize that the responsibility for knowing the name of a class and the responsibility for knowing the name of a message to send to that class may belong in different objects. Just because `Gear` needs to send `diameter` somewhere does not mean that `Gear` should know about `Wheel`.

This leaves the question of where the responsibility for knowing about the actual `Wheel` class lies; the example above conveniently sidesteps this issue, but it is examined in more detail later in this chapter. For now, it's enough to understand that this knowledge does *not* belong in `Gear`.

3.2.2 Isolate Dependencies

It's best to break all unnecessary dependencies but, unfortunately, while this is always *technically* possible, it may not be *actually* possible. When working on an existing application, you may find yourself under severe constraints about how much you can actually change. If prevented from achieving perfection, your goals should switch to improving the overall situation by leaving the code better than you found it.

Therefore, if you cannot remove unnecessary dependencies, you should isolate them within your class. In Chapter 2, you isolated extraneous responsibilities so that they would be easy to recognize and remove when the right impetus came; here you should isolate unnecessary dependencies so that they are easy to spot and reduce when circumstances permit.

Think of every dependency as an alien bacterium that's trying to infect your class. Give your class a vigorous immune system; quarantine each dependency. Dependencies are foreign invaders that represent vulnerabilities, and they should be concise, explicit, and isolated.

Isolate Instance Creation

If you are so constrained that you cannot change the code to inject a `Wheel` into a `Gear`, you should isolate the creation of a new `Wheel` inside the `Gear` class. The intent is to explicitly expose the dependency while reducing its reach into your class.

The next two examples illustrate this idea.

In the first, creation of the new instance of `Wheel` has been moved from `Gear`'s `gear_inches` method to `Gear`'s initialization method. This cleans up the `gear_inches` method and publicly exposes the dependency in the `initialize` method. Notice that this technique unconditionally creates a new `Wheel` each time a new `Gear` is created.

Listing 3.4

```
1 class Gear
2   attr_reader :chainring, :cog, :wheel
3   def initialize(chainring, cog, rim, tire)
4     @chainring = chainring
5     @cog       = cog
6     @wheel     = Wheel.new(rim, tire)
```

```
7   end
8
9   def gear_inches
10    ratio * wheel.diameter
11  end
12  # ...
13  end
14 end
15
16 puts Gear.new(52, 11, 26, 1.5).gear_inches
17 # => 137.0909090909091
```

The next alternative isolates creation of a new `Wheel` in its own explicitly defined `wheel` method. This new method lazily creates a new instance of `Wheel`, using Ruby's `||=` operator. In this case, creation of a new instance of `Wheel` is deferred until `gear_inches` invokes the new `wheel` method.

Listing 3.5

```
1  class Gear
2    attr_reader :chainring, :cog, :rim, :tire
3    def initialize(chainring, cog, rim, tire)
4      @chainring = chainring
5      @cog       = cog
6      @rim       = rim
7      @tire      = tire
8    end
9
10   def gear_inches
11     ratio * wheel.diameter
12   end
13
14   def wheel
15     @wheel ||= Wheel.new(rim, tire)
16   end
17   # ...
18 end
19
20 puts Gear.new(52, 11, 26, 1.5).gear_inches
21 # => 137.0909090909091
```

In both of these examples, Gear still knows far too much; it still takes `rim` and `tire` as initialization arguments, and it still creates its own new instance of `Wheel`. Gear is still stuck to `Wheel`; it can calculate the gear inches of no other kind of object.

However, an improvement *has* been made. These coding styles reduce the number of dependencies in `gear_inches` while publicly exposing Gear's dependency on `Wheel`. They reveal dependencies instead of concealing them, lowering the barriers to reuse and making the code easier to refactor when circumstances allow. This change makes the code more agile; it can more easily adapt to the unknown future.

The way you manage dependencies on external class names has profound effects on your application. If you are mindful of dependencies and develop a habit of routinely injecting them, your classes will naturally be loosely coupled. If you ignore this issue and let the class references fall where they may, your application will be more like a big woven mat than a set of independent objects. An application whose classes are sprinkled with entangled and obscure class name references is unwieldy and inflexible, while one whose class name dependencies are concise, explicit, and isolated can easily adapt to new requirements.

Isolate Vulnerable External Messages

Now that you've isolated references to external class names, it's time to turn your attention to external *messages*, that is, messages that are "sent to someone other than self." For example, the `gear_inches` method below sends `ratio` and `wheel` to `self` but sends `diameter` to `wheel`:

Listing 3.6

```
1 def gear_inches
2   ratio * wheel.diameter
3 end
```

This is a simple method and it contains Gear's only reference to `wheel.diameter`. In this case, the code is fine, but the situation could be more complex. Imagine that calculating `gear_inches` required far more math and that the method looked something like this:

Listing 3.7

```
1 def gear_inches
2   ... a few lines of scary math
3   foo = some_intermediate_result * wheel.diameter
4   ... more lines of scary math
5 end
```

Now `wheel.diameter` is embedded deeply inside a complex method. This complex method depends on `Gear` responding to `wheel` and on `wheel` responding to `diameter`. Embedding this external dependency inside the `gear_inches` method is unnecessary and increases its vulnerability.

Any time you change *anything*, you stand the chance of breaking it; `gear_inches` is now a complex method, and that makes it both more likely to need changing and more susceptible to being damaged when it does. You can reduce your chance of being forced to make a change to `gear_inches` by removing the external dependency and encapsulating it in a method of its own, as in this next example:

Listing 3.8

```
1 def gear_inches
2   #... a few lines of scary math
3   foo = some_intermediate_result * diameter
4   #... more lines of scary math
5 end
6
7 def diameter
8   wheel.diameter
9 end
```

The new `diameter` method is exactly the method that you would have written if you had many references to `wheel.diameter` sprinkled throughout `Gear` and you wanted to DRY them out. The difference here is one of timing; it would normally be defensible to defer creation of the `diameter` method until you had a need to DRY out code; however, in this case, the method is created preemptively to remove the dependency from `gear_inches`.

In the original code, `gear_inches` knew that `wheel` had a `diameter`. This knowledge is a dangerous dependency that couples `gear_inches` to an external object and one of *its* methods. After this change, `gear_inches` is more abstract. `Gear` now isolates `wheel.diameter` in a separate method, and `gear_inches` can depend on a message sent to `self`.

If `Wheel` changes the name or signature of *its* implementation of `diameter`, the side effects to `Gear` will be confined to this one simple wrapping method.

This technique becomes necessary when a class contains embedded references to a *message* that is likely to change. Isolating the reference provides some insurance against being affected by that change. Although not every external method is a candidate for this preemptive isolation, it's worth examining your code, looking for and wrapping the most vulnerable dependencies.

An alternative way to eliminate these side effects is to avoid the problem from the very beginning by reversing the direction of the dependency. This idea will be addressed soon, but first there's one more coding technique to cover.

3.2.3 Remove Argument-Order Dependencies

When you send a message that requires arguments, you, as the sender, cannot avoid having knowledge of those arguments. This dependency is unavoidable. However, passing arguments often involves a second, more subtle dependency. Many method signatures not only require arguments, but they also require that those arguments be passed in a specific, fixed order.

In the following example, `Gear`'s `initialize` method takes three arguments: `chainring`, `cog`, and `wheel`. It provides no defaults; each of these arguments is required. In lines 11–14, when a new instance of `Gear` is created, the three arguments must be passed and they must be passed *in the correct order*.

Listing 3.9

```
1 class Gear
2   attr_reader :chainring, :cog, :wheel
3   def initialize(chainring, cog, wheel)
4     @chainring = chainring
5     @cog       = cog
6     @wheel    = wheel
7   end
8   # ...
9 end
10
11 puts Gear.new(
12     52,
13     11,
14     Wheel.new(26, 1.5)).gear_inches
15 # => 137.0909090909091
```

Senders of `new` depend on the order of the arguments as they are specified in `Gear`'s `initialize` method. If that order changes, all the senders will be forced to change.

Unfortunately, it's quite common to tinker with initialization arguments. Especially early on, when the design is not quite nailed down, you may go through several cycles of adding and removing arguments and defaults. If you use positional arguments, each of these cycles may force changes to many dependents. Even worse, you may find yourself avoiding making changes to the arguments, even when your design calls for them, because you can't bear to change all the dependents yet again.

Use Keyword Arguments

There's a simple way to avoid depending on positional arguments. If you have control over `Gear`'s `initialize` method, change the code to take *keyword* arguments.

The following example illustrates this technique:

Listing 3.10

```
1 class Gear
2   attr_reader :chainring, :cog, :wheel
3   def initialize(chainring:, cog:, wheel:)
4     @chainring = chainring
5     @cog       = cog
6     @wheel     = wheel
7   end
8   # ...
9 end
```

The arguments on line 3 now end in `:`, which denotes that they are keyword arguments. Keyword arguments are referenced just like positional arguments, so lines 4–6 have not changed.

You can pass keyword arguments as a hash, as shown in the following example:

Listing 3.11

```
1 puts Gear.new(
2   :cog      => 11,
3   :chainring => 52,
4   :wheel    => Wheel.new(26, 1.5)).gear_inches
5 # => 137.0909090909091
```

You can also use the explicit keyword syntax:

Listing 3.12

```
1 puts Gear.new(
2   wheel:    Wheel.new(26, 1.5),
3   chainring: 52,
4   cog:      11).gear_inches
5 # => 137.0909090909091
```

Keyword arguments offer several advantages. As you likely noticed in two examples above, keyword arguments may be passed in any order. Additionally, `Gear` is now free to add or remove initialization arguments and defaults, secure in the knowledge that no change will have side effects in other code.

This technique adds verbosity. In many situations verbosity is a detriment, but in this case, it has value. The verbosity exists at the intersection between the needs of the present and the uncertainty of the future. Using positional arguments requires less code today, but you pay for this decrease in volume of code with an increase in the risk that changes will cascade into dependents later.

When Gear switched to keyword arguments, it lost its dependency on argument order but it gained a dependency on the names of the keywords. This change is healthy. The new dependency is more stable than the old, and thus this code faces less risk of being forced to change.

Using keyword arguments requires the sender *and* the receiver of a message to state the keyword names. This results in explicit documentation at both ends of the message. Future maintainers will be grateful for this information.

Keyword arguments are so flexible that the general rule is that you should *prefer* them. While it's certainly true that some argument lists are so stable, and so obvious, that keywords are overkill (for example, what would `Point` take but an `x` and a `y?`), your bias should be toward declaring arguments using keywords. You can always fall back to positional arguments if that technique better suits your specific problem.

Also, it is perfectly acceptable for some classes in your application to take positional arguments and others to take keyword arguments. This is especially true in long-lived applications, where much of the code predates the introduction of keywords. In these cases, as you change or add code, consider using keyword arguments. However, there's no need to proactively retrofit the entire application. Over time, as you touch code, introduce keyword arguments if doing so will add clarity and enable subsequent refactorings.

The remainder of this book uses both types, to supply a flavor of the consequences.

Explicitly Define Defaults

So far, keyword arguments look very similar to hashes. One advantage they have over hashes, however, is that they allow you to set defaults right in the argument list, just like positional arguments. Line 3 below supplies defaults for `chainring` and `cog`.

Listing 3.13

```
1 class Gear
2   attr_reader :chainring, :cog, :wheel
3   def initialize(chainring: 40, cog: 18, wheel:)
4     @chainring = chainring
5     @cog       = cog
6     @wheel    = wheel
7   end
```

```
8 # ...
9 end
10
11 puts Gear.new(wheel: Wheel.new(26, 1.5)).chainring
12 # => 40
```

Notice that the syntax for adding defaults to keyword arguments is a bit different than that of positional arguments. Keywords omit the = operator and state the default directly after the trailing `:`. Adding a default renders the keyword argument optional.

The above syntax is great for supplying simple defaults to optional arguments, but some situations may benefit from a bit more sophistication. For example, line 3 below sets a more complex default by sending a message.

Listing 3.14

```
1 class Gear
2   attr_reader :chainring, :cog, :wheel
3   def initialize(chainring: default_chainring, cog: 18,
4     wheel:)
5     @chainring = chainring
6     @cog       = cog
7     @wheel     = wheel
8   end
9   def default_chainring
10    (100/2) - 10      # silly code, useful example
11  end
12  # ...
13 end
14
15 puts Gear.new(wheel: Wheel.new(26, 1.5)).chainring
16 # => 40
17
18 puts Gear.new(chainring: 52, wheel: Wheel.new(26, 1.5)).
19 chainring
20 # => 52
```

The key to understanding the above code is to recognize that `initialize` executes in the new instance of `Gear`. It is therefore entirely appropriate for `initialize` to send messages to `self`. It's best to embed simple defaults right in the parameter list, but if getting the default requires running a bit of code, don't hesitate to send a message.

Isolate Multiparameter Initialization

So far, all of the examples of removing argument-order dependencies have been for situations where *you* control the signature of the method that needs to change. You will not always have this luxury; sometimes you will be forced to depend on a method that requires positional arguments where you do not own and thus cannot change the method itself.

Imagine that Gear is part of a framework and that its initialization method requires positional arguments. Imagine also that your code has many places where you must create a new instance of Gear. Gear's `initialize` method is *external* to your application; it is part of an interface over which you have no control.

As dire as this situation appears, you are not doomed to accept the dependencies. Just as you would DRY out repetitive code inside of a class, DRY out the creation of new Gear instances by creating a single method to wrap the external interface. The classes in your application should depend on code that you own; use a wrapping method to isolate external dependencies.

In this example, the `SomeFramework::Gear` class is not owned by your application; it is part of an external framework. Its initialization method requires positional arguments. The `GearWrapper` module was created to avoid having multiple dependencies on the order of those arguments. `GearWrapper` isolates all knowledge of the external interface in one place and, equally important, it provides an improved interface for your application.

As you can see in line 22, `GearWrapper` allows your application to create a new instance of `Gear` using keyword arguments.

Listing 3.15

```
1 # When Gear is part of an external interface
2 module SomeFramework
3   class Gear
4     attr_reader :chainring, :cog, :wheel
5     def initialize(chainring, cog, wheel)
6       @chainring = chainring
7       @cog        = cog
8       @wheel      = wheel
9     end
10    # ...
11  end
12 end
13
14 # wrap the interface to protect yourself from changes
15 module GearWrapper
16   def self.gear(chainring:, cog:, wheel:)
```

```
17     SomeFramework::Gear.new(chainring, cog, wheel)
18   end
19 end
20
21 # Now you can create a new Gear using keyword arguments
22 puts GearWrapper.gear(
23   chainring: 52,
24   cog:       11,
25   wheel:     Wheel.new(26, 1.5).gear_inches
26 # => 137.0909090909091
```

There are two things to note about `GearWrapper`. First, it is a Ruby module instead of a class (line 15). `GearWrapper` is responsible for creating new instances of `SomeFramework::Gear`. Using a module here lets you define a separate and distinct object to which you can send the `gear` message (line 22) while simultaneously conveying the idea that you don't expect to have instances of `GearWrapper`. You may already have experience with including modules into classes; in the example above, `GearWrapper` is not meant to be included in another class, it's meant to directly respond to the `gear` message.

The other interesting thing about `GearWrapper` is that its sole purpose is to create instances of some other class. Object-oriented designers have a word for objects like this; they call them *factories*. In some circles, the term *factory* has acquired a negative connotation, but the term as used here is devoid of baggage. An object whose purpose is to create other objects is a *factory*; the word *factory* implies nothing more, and use of it is the most expedient way to communicate this idea.

The above technique for replacing positional arguments with keywords is perfect for cases where you are forced to depend on external interfaces that you cannot change. Do not allow these kinds of external dependencies to permeate your code; protect yourself by wrapping each in a method that is owned by your own application.

3.3 Managing Dependency Direction

Dependencies always have a direction; earlier in this chapter it was suggested that one way to manage them is to reverse that direction. This section delves more deeply into how to decide on the direction of dependencies.

3.3.1 Reversing Dependencies

Every example used thus far shows `Gear` depending on `Wheel` or `diameter`, but the code could easily have been written with the direction of the dependencies reversed.

Wheel could instead depend on Gear or ratio. The following example illustrates one possible form of the reversal. Here Wheel has been changed to depend on Gear and gear_inches. Gear is still responsible for the actual calculation, but it expects a diameter argument to be passed in by the caller (line 8).

Listing 3.16

```
1 class Gear
2   attr_reader :chainring, :cog
3   def initialize(chainring:, cog:)
4     @chainring = chainring
5     @cog       = cog
6   end
7
8   def gear_inches(diameter)
9     ratio * diameter
10  end
11
12  def ratio
13    chainring / cog.to_f
14  end
15  # ...
16 end
17
18 class Wheel
19   attr_reader :rim, :tire, :gear
20   def initialize(rim:, tire:, chainring:, cog:)
21     @rim = rim
22     @tire = tire
23     @gear = Gear.new(chainring: chainring, cog: cog)
24   end
25
26   def diameter
27     rim + (tire * 2)
28   end
29
30   def gear_inches
31     gear.gear_inches(diameter)
32   end
33   # ...
34 end
35
36 puts Wheel.new(
```

```
37         rim:      26,  
38         tire:     1.5,  
39         chainring: 52,  
40         cog:      11).gear_inches  
41 # => 137.0909090909091
```

This reversal of dependencies does no apparent harm. Calculating `gear_inches` still requires collaboration between `Gear` and `Wheel` and the result of the calculation is unaffected by the reversal. One could infer that the direction of the dependency does not matter, that it makes no difference whether `Gear` depends on `Wheel` or vice versa.

Indeed, in an application that never changed, your choice would not matter. However, your application *will* change, and it's in that dynamic future where this present decision has repercussions. The choices you make about the direction of dependencies have far-reaching consequences that manifest themselves for the life of your application. If you get this right, your application will be pleasant to work on and easy to maintain. If you get it wrong, then the dependencies will gradually take over and the application will become harder and harder to change.

3.3.2 Choosing Dependency Direction

Pretend for a moment that your classes are people. If you were to give them advice about how to behave, you would tell them to *depend on things that change less often than you do*.

This short statement belies the sophistication of the idea, which is based on three simple truths about code:

- Some classes are more likely than others to have changes in requirements.
- Concrete classes are more likely to change than abstract classes.
- Changing a class that has many dependents will result in widespread consequences.

There are ways in which these truths intersect, but each is a separate and distinct notion.

Understanding Likelihood of Change

The idea that some classes are more likely to change than others applies not only to the code that you write for your own application but also to the code that you use but did *not* write. The Ruby base classes and the other framework code that you rely on both have their own inherent likelihood of change.

You are fortunate in that Ruby base classes change a great deal less often than your own code. This makes it perfectly reasonable to depend on the `*` method, as `gear_inches` quietly does, or to expect that Ruby classes `String` and `Array` will continue to work as they always have. Ruby base classes always change less often than your own classes, and you can continue to depend on them without another thought.

Framework classes are another story; only you can assess how mature your frameworks are. In general, any framework you use will be more stable than the code you write, but it's certainly possible to choose a framework that is undergoing such rapid development that its code changes more often than yours.

Regardless of its origin, every class used in your application can be ranked along a scale of how likely it is to undergo a change relative to all other classes. This ranking is one key piece of information to consider when choosing the direction of dependencies.

Recognizing Concretions and Abstractions

The second idea concerns itself with the concreteness and abstractness of code. The term *abstract* is used here just as Merriam-Webster defines it, as “disassociated from any specific instance,” and, as so many things in Ruby, represents an idea about code as opposed to a specific technical restriction.

This concept was illustrated earlier in the chapter during the section on injecting dependencies. There, when `Gear` depended on `Wheel` and on `Wheel.new` and on `Wheel.new(rim, tire)`, it depended on extremely concrete code. After the code was altered to inject a `Wheel` into `Gear`, `Gear` suddenly began to depend on something far more abstract, that is, the fact that it had access to an object that could respond to the `diameter` message.

Your familiarity with Ruby may lead you to take this transition for granted, but consider for a moment what would have been required to accomplish this same trick in a statically typed language. Because statically typed languages have compilers that act like unit tests for types, you would not be able to inject just any random object into `Gear`. Instead you would have to declare an *interface*, define `diameter` as part of that interface, include the interface in the `Wheel` class, and tell `Gear` that the class you are injecting is a *kind of* that interface.

Rubists are justifiably grateful to avoid these gyrations, but languages that force you to be explicit about this transition do offer a benefit. They make it inescapably and explicitly clear that you are defining an abstract interface. It is impossible to create an abstraction unknowingly or by accident; in statically typed languages, defining an interface is *always* intentional.

In Ruby, when you inject `Wheel` into `Gear` such that `Gear` then depends on a *Duck* who responds to `diameter`, you are, however casually, defining an interface. This interface is an abstraction of the idea that a certain category of things will have a `diameter`. The abstraction was harvested from a concrete class; the idea is now “disassociated from any specific instance.”

The wonderful thing about abstractions is that they represent common, stable qualities. They are less likely to change than are the concrete classes from which they were extracted. Depending on an abstraction is always safer than depending on a concretion because by its very nature, the abstraction is more stable. Ruby does not make you explicitly declare the abstraction in order to define the interface, but for design purposes, you can behave as if your virtual interface is as real as a class. Indeed, in the rest of this discussion, the term *class* stands for both *class* and this kind of *interface*. These interfaces can have dependents and so must be taken into account during design.

Avoiding Dependent-Laden Classes

The final idea, the notion that having dependent-laden objects has many consequences, also bears deeper examination. The consequences of changing a dependent-laden class are quite obvious—not so apparent are the consequences of even *having* a dependent-laden class. A class that, if changed, will cause changes to ripple through the application will be under enormous pressure to *never* change. Ever. Under any circumstances whatsoever. Your application may be permanently handicapped by your reluctance to pay the price required to make a change to this class.

Finding the Dependencies That Matter

Imagine each of these truths as a continuum along which all application code falls. Classes vary in their likelihood of change, their level of abstraction, and their number of dependents. Each quality matters, but the interesting design decisions occur at the place where *likelihood of change* intersects with *number of dependents*. Some of the possible combinations are healthy for your application; others are deadly.

Figure 3.2 summarizes the possibilities.

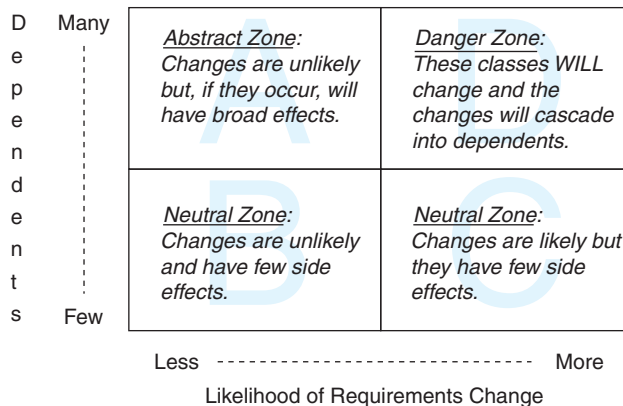


Figure 3.2 Likelihood of change versus number of dependents

The likelihood of requirements change is represented on the horizontal axis. The number of dependents is on the vertical. The grid is divided into four zones, labeled A through D. If you evaluate all of the classes in a well-designed application and place them on this grid, they will cluster in Zones A, B, and C.

Classes that have little likelihood of change but contain many dependents fall into Zone A. This zone usually contains abstract classes or interfaces. In a thoughtfully designed application, this arrangement is inevitable; dependencies cluster around abstractions because abstractions are less likely to change.

Notice that classes do not become abstract because they are in Zone A; instead they wind up here precisely because they are *already* abstract. Their abstract nature makes them more stable and allows them to safely acquire many dependents. While residence in Zone A does not guarantee that a class is abstract, it certainly suggests that it ought to be.

Skipping Zone B for a moment, Zone C is the opposite of Zone A. Zone C contains code that is quite likely to change but has few dependents. These classes tend to be more concrete, which makes them more likely to change, but this doesn't matter because few other classes depend on them.

Zone B classes are of the least concern during design because they are almost neutral in their potential future effects. They rarely change and have few dependents.

Zones A, B, and C are legitimate places for code; Zone D, however, is aptly named the Danger Zone. A class ends up in Zone D when it is guaranteed to change *and* has many dependents. Changes to Zone D classes are costly; simple requests become coding nightmares as the effects of every change cascade through each dependent. If you have a very specific *concrete* class that has many dependents and you believe it resides in Zone A, that is, you believe it is unlikely to change, think again. When a concrete class has many dependents, your alarm bells should be ringing. That class might actually be an occupant of Zone D.

Zone D classes represent a danger to the future health of the application. These are the classes that make an application painful to change. When a simple change has cascading effects that force many other changes, a Zone D class is at the root of the problem. When a change breaks some far away and seemingly unrelated bit of code, the design flaw originated here.

As depressing as this is, there is actually a way to make things worse. You can guarantee that any application will gradually become unmaintainable by making its Zone D classes *more* likely to change than their dependents. This maximizes the consequences of every change.

Fortunately, understanding this fundamental issue allows you to take preemptive action to avoid the problem.

Depend on things that change less often than you do is a heuristic that stands in for all the ideas in this section. The zones are a useful way to organize your thoughts, but in the fog of development, it may not be obvious which classes go where. Very

often you are exploring your way to a design, and at any given moment the future is unclear. Following this simple rule of thumb at every opportunity will cause your application to evolve a healthy design.

3.4 Summary

Dependency management is core to creating future-proof applications. Injecting dependencies creates loosely coupled objects that can be reused in novel ways. Isolating dependencies allows objects to quickly adapt to unexpected changes. Depending on abstractions decreases the likelihood of facing these changes.

The key to managing dependencies is to control their direction. The road to maintenance nirvana is paved with classes that depend on things that change less often than they do.

This page intentionally left blank

Index

Unspecified

`||=` operator, 45

A

Abstract

classes, 118, 123–125, 242–244
definition of, 56
documentation, supplying, 195
superclass, creating, 117–120

Abstractions

extracting, 150–153
insisting on, writing inheritable code, 158–159
recognizing, 56–57
supporting, in intentional testing, 195–196

Abstractions, finding

abstract superclass, creating, 117–120
overview of, 116–117
promoting abstract behavior, 120–123
separating from concretions, 123–125
template method pattern, 125–129

Across-class types, 86

Ad hoc methods, 156

Ad infinitum, 86

Aggregation, composition vs., 184–185

Agile, 7–9

Antipatterns

recognizing, 158
understanding, 110–111

Argument-order dependencies, removing

keyword arguments, using, 49–50
multiparameter initialization, isolation, 52–53
overview of, 48, 50–51

Attributes, OO language

and, 12

Automatic message delegation, 105–106

B

behaves-like-a relationships, duck types for, 190

Behavior

acquiring via inheritance.

See Inheritance

data structures, hiding, 27–29

depending on, instead of data, 24–29

duck types and. *See* Duck types

instance variables, hiding, 24–27

mock tests of, 219–221

promoting abstract, 120–123
sharing role. *See* Modules

single responsibility and, 21
subclass, 237–239
testing inherited code for unique, 240–244

Behavior-Driven Development (BDD), 200–201

Big Up Front Design (BUFD), 8

Booch, Grady, 189

Break-even point, for design, 10

Brittleness, test

solving problem of, 227–233
stubbing and mocking, creating, 213

Bugs, finding, 195

C

Case statements that switch on class
hidden ducks, recognizing, 95–96

`is_a?` 96–97

`kind_of?` 96–97

overview of, 95–96

`responds_to?` 97

Categories, testing with object, 201

category variable

antipatterns, recognizing, 158
embedded types, finding, 111–112

- CBO (coupling between objects), 39–40
- Change
 - Agile, guaranteeing, 9
 - code, organizing to allow for, 16–17
 - designing to reduce cost of, 3–4
 - managing dependencies. *See* Dependencies
 - in OO languages, 3
 - as unavoidable and inevitable, 2
 - writing code to embrace, 24–33
- Class-based OO languages, 12–13
- Class class, 12
- Classes. *See also* Single responsibility, classes with
 - abstract, 118, 123–125
 - avoiding dependency-laden, 57
 - case statements that switch on, 95–97
 - code, organizing to allow for changes, 16–17
 - concrete, 107–109, 123–125, 241–242
 - deciding what belongs in, 16–17
 - decoupling, writing inheritable code, 21
 - dependency injection using, 208–210
 - dependent-laden, avoiding, 57
 - grouping methods into, 16
 - predefined, 12
 - references to. *See* Loosely coupled code, writing
 - reusable, 21
 - Ruby-based vs. framework, 55–56
 - type and category used in, 111–112
 - virtual, 63
- Classical inheritance, 105–106, 185
- Code
 - concrete, writing, 147–150
 - dependence on behavior, not data, 24–29
 - dependency injection to shape, 41–44
 - embracing change, writing, 16–17, 24–33, 193
 - error messages for failed, 129
 - inheritable, writing, 158–161
 - initialization, 120–121
 - interface, writing, 77–80
 - Law of Demeter (LoD) rules for, 80–84
 - loosely coupled, writing, 41–48
 - open-closed, 186
 - putting into modules, 143
 - single responsibility, enforcing, 29–33
 - that relies on ducks, writing, 95–100
 - truths about, 55
- Code arrangement technique, classical inheritance as, 185
- Code, testing inherited
 - inherited interface, specifying, 233–236
 - subclass responsibilities, specifying, 236–240
 - unique behavior, testing, 240–244
- Code-to-code over-coupling, 41
- Cohesion, 22
- Command messages, 197, 218–221
- Comments, 31
- Communication patterns, 61–63
- Compile-time type checking, 102–103
- Composition
 - aggregation as special form of, 184–185
 - benefits of, 188
 - of bicycle, 181–184
 - of bicycle of parts, 163–168
 - combining objects with, 163
 - consequences of, accepting, 186–189
 - costs of, 188–189
 - deciding between inheritance and, 185–191
 - defined, 163
 - for has-a relationships, 163, 190–191
 - manufacturing parts, 176–181
 - of parts object, 168–176
- Concrete classes
 - inheritance and, 107–109
 - likelihood of change in, 55, 58
 - separating abstract from, 123–125
 - testing concrete subclass behavior, 241–242
- Concretions
 - abstractions separated from, 123–125
 - duck typing and, 91, 94–95
 - inheritance and, 107–109
 - recognizing, 56–57
 - writing, 147–150
- config array, 176–179
- Context
 - independence, seeking, 72–74
 - minimizing in public interfaces, 80
- Contract, honoring, 159, 161
- Costs
 - of composition, 188–189
 - of inheritance, 187–188
 - reducing with duck typing. *See* Duck typing
 - of testing. *See* Testing
- Coupling
 - decoupling classes, writing inheritable code, 160
 - decoupling subclasses with hook messages, 134–139
 - knowing what to test, 197
 - loosely coupled code. *See* Loosely coupled code, writing

- overview of, 129
- sharing code via modules and, 151
- between superclasses and subclasses, 129–134
- tests on code, and tight, 41
- Coupling between objects (CBO), 39–40
- D**
- Data
 - depending on behavior instead of, 24–29
 - instance variables, hiding, 24–27
 - in object-oriented languages, 11–13
 - in procedural languages, 11 structures, hiding, 27–29 types, 11–12
- Decoupling
 - classes, writing inheritable code, 159
 - subclasses, with hook messages, 134–139
- Defaults, explicitly defining, 50–51
- Delegation
 - automatic message, in inheritance, 105–106
 - automatic message, in modules, 143
 - composition and, 184
- Demeter. *See* Law of Demeter (LoD)
- Dependencies
 - coupling between objects, 39–40
 - delegation creating, 184
 - deleting unused interfaces, 204
 - direction of. *See* Dependency direction
 - injecting. *See* Dependency injection
 - interfaces and, 64
 - isolating, 44–48
 - loosely coupled code. *See* Loosely coupled code, writing
 - objects speaking for themselves, 147
 - other, 40–41
 - overview of, 37
 - recognizing, 39
 - removing argument-order, 48–53
 - removing unnecessary, 146–147
 - reversing, 53–55
 - scheduling duck type, discovering, 146–147
 - summary, 59
 - tolerating change by, 3
 - understanding, 38–39
- Dependency direction
 - abstractions, recognizing, 56–57
 - change in, likelihood of, 55–56
 - concretions, recognizing, 56–57
 - dependent-laden classes, avoiding, 57
 - finding dependencies that matter, 57–59
 - overview of, 55
 - reversing, 53–55
- Dependency injection
 - failure of, 62
 - in loosely coupled code, 41–44
 - as roles, 210–211
 - to shape code, 43–44
 - using classes, 208–210
- Dependency Inversion
 - principle, 4
- Dependent-laden classes, avoiding, 57
- Design
 - act of, 6–10
 - definition of, 3–4
 - failure in, 6–7
 - judging, 9–10
 - patterns, 4–5
 - principles, 4–5
 - problems solved by, 2–3
 - when to, 7–9
 - why change is hard, 3
- Design decisions
 - deferring, 195
 - when to make, 22–23
- Design flaws, exposing, 196
- Design patterns, 6
- Design Patterns: Elements of Reusable Object-Oriented Software (Gamma, Helm, Johnson and Vlissides), 6, 189
- Design principles, 4–5, 197
- Design tools, 4–6
- Documentation
 - of duck types, 98
 - of error messages, in template method, 129
 - of roles, testing used in, 214–215
 - supplying, in testing, 195
- Domain objects, 66, 95, 201
- Doubles
 - creating test, 211–214
 - role tests to validate, 227–233
- DRY (Don't Repeat Yourself), 4–5, 24, 28–29, 47, 52, 198
- Duck types
 - choosing wisely, 98–100
 - defined, 85
 - documenting, 98
 - finding, 90–94
 - hidden, recognizing, 95–97
 - overlooking, 86–88
 - sharing code between, 98
 - testing roles, 221–227
 - testing roles to validate doubles, 227–233
 - trust in, placing, 97–98
- Duck typing
 - for behaves-like-a relationships, 190
 - code that relies on, writing, 95–100
 - consequences of, 94–95
 - dynamic typing and, 100–103
 - fear of, conquering, 100–103
 - finding roles, 142–143
 - inheritable code, writing for, 158
 - overview of, 85
 - problem, compounding, 88–90
 - scheduling, discovering, 146–147

- sharing code via Ruby
 - modules, 112
 - static typing and,
 - 100–101
 - summary, 103
 - understanding, 85–94
- Duplication, removing test, 196
- Dynamic typing, 100–103
- E**
- Embedded types of
 - inheritance
 - finding, 111–112
 - multiple, 109–111
- Error messages, code that fails
 - with reasonable, 129
- Explicit interfaces, creating, 77–78
- F**
- Factories, 53, 176–179
- Failures, design, 6–7
- Family tree image, inheritance, 112
- Feathers, Michael, 4
- Fixed order arguments, 48–53
- Fowler, Martin, 193
- Framework, choosing testing, 200
- Framework class, 56
- G**
- Gamma, Erich, 6, 189
- Gang of Four (GoF), 6
- Gear inches, 20–21
- H**
- has-a relationships
 - using composition for, 163, 184–185, 190–191
 - vs. is-a relationships, 191
- Hashes, for initialization
 - arguments, 49–50
- Helm, Richard, 6, 189
- Highly cohesive class, 22
- Hook messages, 134–140
- How to test, knowing, 200–202
- “How” vs. “what,” 70–72
- Hunt, Andy, 4
- I**
- Incoming messages, testing
 - documenting roles via, 214–215
- injecting dependencies as
 - roles, 210–211
- injecting dependencies using
 - classes, 208–210
- interfaces, deleting unused, 204
- isolating object under test, 206–208
- knowing what to test, 197–198
- overview of, 202–203
- proving public interface, 204–206
- Inheritable code, writing
 - abstraction, insisting on, 158–159
 - antipatterns, recognizing, 158
 - contract, honoring, 159
 - decouple classes,
 - preemptively, 160
 - shallow hierarchies, creating, 160–161
 - template method pattern,
 - using, 160
- Inheritance
 - behavior acquired through, 105
 - benefits of, 186
 - choosing, 112–114
 - classical, 105–106
 - composition and, deciding
 - between, 185–191
 - concretions and, 107–109
 - costs of, 187–188
 - coupling and, 129–134
 - decoupling subclasses with
 - hook messages, 134–139
 - drawing inheritance
 - relationships, 114
 - embedded types of, 109–112
 - family tree image of, 112
 - implying, 117
 - inherited code, testing, 233–244
 - for is-a relationships, 186, 189–190
 - misapplying, 114–116
 - multiple, 112
 - overview of, 106–107
 - polymorphism, achieving
 - via, 95
 - problem solved by, 112
 - recognizing where to use, 107–114
 - of role behavior, 157–158
 - single, 112
 - subclasses/superclasses,
 - coupling between, 129–139
 - summary, 139–140
 - using for is-a relationships, 186, 189–190
- Inheritance, finding
 - abstraction
 - creating abstract from
 - concrete, 123–125
 - creating abstract superclass, 117–120
 - implementing every template
 - method, 127–129
 - overview of, 116–117
 - promoting abstract behavior, 120–123
 - using the template method
 - pattern, 125–127
- Inherited code, testing
 - inherited interface,
 - specifying, 233–236
 - subclass responsibilities,
 - specifying, 236–240
 - unique behavior, 240–244
- Inherited interface, specifying, 233–236
- Initialization arguments
 - dependency injection and, 51
 - isolating instance creation, 46
 - keyword arguments
 - using, 49
 - manufacturing parts, 179–180
 - removing argument-order
 - dependencies, 48
- Injection. *See* Dependency injection

- Instance creation, isolating, 44–45
- Instance variables, hiding, 24–27
- Intention, constructing, 65–66
- Intentional testing
 - knowing how to test, 200–202
 - knowing what to test, 196–199
 - knowing when to test, 199–200
 - knowing your intentions, 194–196
 - overview of, 194
- Interface Segregation principle, 4
- Interfaces. *See also* Private interfaces; Public interfaces
 - code, writing for, 77–80
 - deleting unused, 204
 - dependencies and, 64–65
 - inherited interface, specifying, 233–236
 - Law of Demeter and, 80–84
 - overview of, 61
 - responsibilities and, 64–65
 - summary, 84
 - understanding, 61–63
 - wrapping dependencies in external, 52–53
- is-a relationships
 - using inheritance for, 186, 189–190
 - vs. has-a relationships, 191
- is_a? 96–97
- Isolation
 - of dependencies, 44–48
 - of external messages, 46–48
 - of instance creation, 44–46
 - of multiparameter initialization, 52–53
 - of object under test, 206–208
 - of responsibilities in classes, 32–33
- Iterations, Agile development, 7–8
- J**
- Java, 102, 118
- JavaScript, 106
- Johnson, Ralph, 6, 189
- Judging design, 9–10
- K**
- Keyword arguments
 - adding defaults to, 50–51
 - argument-order dependencies, removing, 49–50
 - multiparameter initialization, isolating, 52–53
- Keywords, in Ruby, 78–79
- kind_of? 96–97
- L**
- Languages
 - object-oriented, 11–13
 - OO design failure and, 7
 - procedural, 11
- Law of Demeter (LoD)
 - design principle, 4–5
 - listening to, 83–84
 - overview of, 80
 - violations, 40–41, 81–83
- Likelihood of change
 - dependencies that matter and, 57–59
 - in embedded references to message, 47–48
 - understanding, 55–56
- Liskov, Barbara, 159
- Liskov Substitution Principle (LSP), 4, 159, 161, 234–235
- Lookups, method, 153–157
- Loosely coupled code, writing
 - inject dependencies, 41–44
 - isolate dependencies, 44–48
 - overview of, 41
 - remove argument-order dependencies, 48–53
- M**
- Managing dependencies. *See* Dependencies
- Martin, Robert, 4
- Message chaining, 40–41, 81–83
- Messages. *See also* Incoming messages, testing applications, creating based on, 77 asking “what” vs. telling “how” and, 70–72 automatic message delegation, 105–106 command, proving, 218–221 defining, 63 designing based on, 69–70 external, isolating vulnerable, 46 forwarding, via classical inheritance, 112 as foundation of OO system, 15 interface defined by, 63 isolating vulnerable external, 46–48 Law of Demeter (LoD) and, 80–84 likely to change, embedded references to, 47–48 method lookup for, 154–157 objects discovered by, 75–77 outgoing, testing, 217–221 polymorphism and, 94–95 query, ignoring, 217–218 sequence diagrams, using for, 66–70
- Methods
 - defining in module, 143
 - extracting extra responsibilities from, 29–32
 - grouping into classes, 16
 - looking up, 153–157, 161
 - private, testing, 216–217
 - wrapper, 24–27
- Metrics, 9–10, 110
- Meyer, Bertrand, 189
- Minitest framework, 200, 201–202
- Mock tests
 - brittleness and, 213
 - proving command messages, 219–221
 - proving outgoing messages, 226–227
 - solving brittleness problem, 227–233

- Modules
 - defined, 143
 - inheritable code, writing, 158–161
 - method lookup for, 155–157
 - polymorphism and, 94–95
 - sharing code via, 141
 - summary, 161
 - understanding roles. *See* Roles
- Monkey patching, 99
- Multiparameter initialization, isolating, 52–53
- Multiple inheritance, 112
- N**
- Names, class
 - dependency injection and, 41–44
 - isolating references to external, 46
 - rules of inheritance, 116–117
- NASA Goddard Space Flight Center applications, 5
- `nil`, 99, 103, 113
- `NilClass`, 99, 113
- O**
- Object-Oriented Analysis and Design (Booch), 189
- Object-oriented design (OOD)
 - brief introduction to, 10
 - how it fails, 6–7
 - judging, 9–10
 - object-oriented languages, 11–12
 - overview of, 1–2
 - patterns, 4–5
 - praise of, 2–4
 - principles, 4–5
 - procedural languages, 11
 - summary, 13
- Object-oriented languages, 11–13
- Objects
 - combined with composition. *See* Composition
 - discovering using messages, 75–77
 - duck typing and, 85–86
 - messages used to discover, 75–77
 - sequence diagrams for, 66–70
 - speaking for themselves, 147
 - trusting other, 70–72
- Open-closed code, 186
- Open-Closed principle, 4, 186
- `OpenStruct` class, 179–181, 182–184
- Outgoing messages, testing
 - ignoring query messages, 217–218
 - knowing what to test, 197–198
 - overview of, 217
 - proving command messages, 218–221
- P**
- Parts object
 - composed bicycle, 181–185
 - composing bicycle of parts and, 163–164
 - creating, 168–173
 - creating `PartsFactory`, 176–179
 - hierarchy, creating, 165–168
 - leveraging `PartsFactory`, 179–181
 - making more like array, 172–176
 - manufacturing, 176–181
 - updating `Bicycle` class, 164–165
- Patterns, interface message, 61–63
- Point-of-view, in testing, 201
- Polymorphism, 94–95
- Positional arguments, 49–50, 52–53
- Private interfaces
 - creating explicit, 78
 - defined, 64
 - depending on, caution, 79–80
 - public vs., 64–65
 - private keyword, 78–79
- Private methods, testing
 - choosing, 216–217
 - ignoring, 216
- removing from class under test, 216
- Procedural languages, 11
- Promotion failure, 120–123
- Public interfaces
 - code, writing for, 77–80
 - context independence, seeking, 72–74
 - context, minimizing in, 80
 - defining, 64, 65
 - Demeter violations in, 83–84
 - duck types as. *See* Duck types
 - example application, bicycle touring company, 65
 - intention, constructing, 65–66
 - knowing what to test and, 197–199
 - message-based application, creating, 77
 - messages used to discover other objects, 75–77
 - of others, honoring, 79–80
 - proving, 204–206
 - sequence diagrams, using, 66–70
 - testing incoming messages, 203
 - trusting other objects, 74–75
 - understanding, 62
 - “what” vs “how,” importance of, 70–72
- Q**
- Query messages, 198, 217–218
- R**
- RDD (Responsibility-Driven Design), 22
- Refactoring
 - barriers, reducing to, 217
 - in extracting extra responsibilities from methods, 30–32
 - intentional testing, and, 197
 - keyword arguments and, 50
 - rule for, 123
 - strategies, deciding between, 122–123

- testing roles and, 223–224, 229
- in writing changeable code, 193–194
- Refactoring: Improving the Design of Existing Code (Fowler), 193
- Relationships
 - aggregation and, 184–185
 - choosing, 189
 - drawing inheritance, 114
 - inheritance and, 105–106, 112–114
 - use composition for has-a, 190–191
 - use duck types for
 - behaves-like-a, 190
 - use inheritance for is-a, 189–190
- responds_to? 97
- Responsibilities
 - designing interfaces, 64
 - organizing, 143–145
- Responsibility-Driven Design (RDD), 22
- Reusability
 - single responsibility classes
 - for, 21
 - single responsibility methods
 - for, 32
 - test-first and, 199–200
- Reversal of dependencies, 53–55
- Roles
 - concrete code, writing, 147–150
 - dependencies, removing unnecessary, 146–147
 - extracting abstraction and, 150–152
 - finding, 142–143
 - inheritable code, writing, 158–161
 - inheriting behavior of, 157–158
 - injecting dependencies as, 210–211
 - test doubles, creating to play, 211–214
 - testing, in duck typing, 219–221
 - testing to document, 214–215
 - tests to validate doubles, 227–233
 - understanding, 142
 - using duck types for
 - behaves-like-a relationships, 190
- RSpec framework, 200
- Ruby, 7, 9–13
- Ruby-based class vs. framework, 55–56
- S**
- Sequence diagrams
 - duck typing and, 90–92, 146–147
 - message-based applications,
 - creating, 77
 - using, 66–70
- Shallow hierarchies,
 - inheritable code, 95
- Single inheritance, 112
- Single responsibility, classes
 - with
 - code embracing change,
 - writing, 24–29
 - creating, 17–23
 - deciding what belongs in class, 16–17
 - design decisions, when to make, 22–23
 - determining, 22
 - discovering objects using messages, 75–77
 - enforcing everywhere, 29–32
 - example application,
 - bicycles and gears, 17–21
 - extra responsibilities,
 - isolating, 32–33
 - overview of, 15–16
 - real wheel, 33–35
 - summary, 35
 - why it matters, 21
- Single Responsibility Principle (SRP), 22, 186
- SLOC (source lines of code), 9–10
- SOLID (Single Responsibility, Open-Closed, Liskov Substitution, Interface Segregation, and Dependency Inversion)
 - design principles, 4–5, 159
- Source code repository, 61
- Source lines of code (SLOC), 9–10
- Specializations of their superclasses, subclasses
 - as, 113, 116, 118
- Spike a problem, 200
- Static typing
 - duck types, subverting with, 100–101
 - vs. dynamic typing, 102–103
- String class, 12
- String data type, 11–12
- String objects, 11–12
- Struct class
 - adding new methods, 32–35
 - OpenStruct class vs., 179–181
 - wrapping data structures, 28–29
- Stubbing
 - brittleness caused by, 213
 - solving brittleness problem, 227–233
- Styles, testing, 200–201
- Subclasses
 - concrete behavior of, testing, 241–242
 - coupling between superclasses and, 129–134
 - decoupling using hook messages, 134–139
 - inheritable code, writing, 158–159
 - inheritance relationships and, 106
 - responsibilities of,
 - specifying, 236–240
 - rules of inheritance, 117
 - shared behavior across set of, 118
 - in single inheritance, 112

- as specializations of their superclasses, 113, 117
- template method pattern and, 125–127
- Superclasses
 - abstract behavior, promoting to, 120–123
 - abstract, creating, 117–120
 - automatic delegation of messages to, 106, 113
 - confirming enforcement of, 239–240
 - coupling between subclasses and, 129–134
 - inheritable code, writing, 158–159
 - looking up methods, 154–157
 - misapplying inheritance, 114–116
 - rules of inheritance, 117
 - single inheritance and, 112
 - subclasses as specializations of, 113, 117
 - template method pattern and, 125–127
- T**
- Technical debt, 10, 80
- Template method pattern
 - coupling between subclasses/superclasses, 131–134
 - decoupling subclasses with hook messages, 138–139
 - defined, 125
 - error messages for failed code, 127–129
 - inheritable code, writing, 159
 - sharing code via modules and, 151–153
 - summary, 140
 - using, 125–127
- Test-Driven Development (TDD), 200–201
- Test-to-code over-coupling, 41
- Testing
 - abstractions, supporting, 195–196
 - bugs, finding, 195
 - cost-effective, designing, 193–244
 - design decisions, deferring, 195
 - design flaws, exposing, 196
 - duck types, 221–233
 - incoming messages, 202–215
 - inherited code, 233–244
 - intentional testing, 194–202
 - knowing how to test, 200–202
 - knowing what to test, 196–199
 - knowing when to test, 199–200
 - knowing your intentions, 194–196
 - outgoing messages, 217–221
 - over-coupling, 41
 - overview of, 193–194
 - summary, 244
 - test doubles, creating, 211–214
- Thomas, Dave, 4
- Time interval that matters, 10
- Touch of Class: Learning to Program Well with Objects and Contracts (Meyer), 189
- TRUE (Transparent, Reasonable, Usable, and Exemplary) code, 16–17, 23
- Trust
 - of other objects, 70–72
 - of your ducks, 97–98
- Type declarations, dynamic typing and, 102
- Type variable, 111–112, 158
- Types. *See also* Duck typing
 - embedding multiple, 109–111
 - finding embedded, 111–112
- U**
- Unified Modeling Language (UML) class diagrams, 66–70, 114
- Updating class, for composition, 164–165
- V**
- Variables, defining, 11
- Violations, Law of Demeter (LoD), 40–41, 81–83
- Virtual class, 63
- Vlissides, John, 6, 189
- W**
- What to test, knowing, 196–199
- “What” vs. “how,” importance of, 70–72
- When to test, knowing, 199–200
- Wilkerson, Brian, 22
- Wirfs-Brock, Rebecca, 22
- Within-class types, 86
- Wrapper method, 25, 82

Credits

- Chapter 1, “a way to produce cheaper and higher quality software”: Laing, Victor & Coleman, Charles. (2001). Principal Components of Orthogonal Object-Oriented Metrics (323-08-14).
- Chapter 1, “simple and elegant solutions to specific problems in object-oriented software design”, “make your own designs more flexible, modular, reusable and understandable”: Gamma, E., et al. (1995). Design Patterns: Elements of Reusable Object-Oriented Software. Reading, MA: Addison-Wesley.
- Chapter 2, “A class has responsibilities that fulfill its purpose”: Rebecca Wirfs-Brock; Brian Wilkerson’s idea, Responsibility-Driven Design (RDD).
- Chapter 2, “a convenient way to bundle a number of attributes together, using accessor methods, without having to write an explicit class.”: <http://ruby-doc.org/core/classes/Struct.html>
- Chapter 3, abstract: “disassociated from any specific instance”: Merriam-Webster.
- Chapter 8, “Inheritance is specialization”: Bertrand Meyer, Touch of Class: Learning to Program Well with Objects and Contracts.
- Chapter 8, “Use composition when...sum of its parts” Booch, G., et al. (2007). Object-Oriented Analysis and Design with Applications. Upper Saddle River: Addison-Wesley.
- Chapter 8, “Inheritance is best...small amounts of new code”: Gamma, E., et al. (1995). Design Patterns: Elements of Reusable Object-Oriented Software. Reading, MA: Addison-Wesley.
- Chapter 9, “Refactoring is the process...improves the internal structure”: Fowler, M., et al. (1999). Refactoring: Improving the Design of Existing Code. Reading, MA: Addison-Wesley.
- Cover: Wheels of Progress series. Backdrop of gears and fractal radial elements on the subject of science, technology and education. Agsandrew/Shutterstock.