

JULIEN DANJOU

SERIOUS PYTHON

BLACK BELT ADVICE



Acknowledgements

Writing this first book has been a tremendous effort. Looking back, I had no clue how crazy this journey would be but also now idea how fulfilling it has been.

They say that if you want to go fast you should go alone, but that if you want to go far you should go together. As this is the fourth edition of the original book I wrote, I would not have made it without people along the way. This is a team effort and I would like to thank everyone who participated.

Most of the interviewees gave me their time and trust without a second thought, and I owe them a lot of what we teach in this book: Doug Hellmann for its great advice about building libraries, Joshua Harlow for his good mood and knowledge about distributed systems, Christophe de Vienne for his experience around building a framework, Victor Stinner for his incredible knowledge about CPython, Dimitri Fontaine for his wisdom on databases, Robert Collins for messing up with testing, Nick Coghlan for his work in getting Python in a better shape and Paul Tagliamonte for his amazing hacker spirit.

Thanks to the No Starch crew for working with me on bringing this book to a brand new level — especially to Liz Chadwick for her editing skills, Laurel Chun for keeping me on track and Michael Driscoll for his technical hindsight.

My gratitude also goes to the Free Software communities who help me grow and gave me the willingness to share my knowledge back, especially to the Python community which always has been welcoming and enthusiastic.

Contents

1	Starting Your Project	21
1.1	Versions of Python	21
1.2	Laying Out Your Project	23
1.2.1	What to Do	23
1.2.2	What Not to Do	25
1.3	Version Numbering	26
1.4	Coding Style & Automated Checks	28
1.4.1	Tools to Catch Style Errors	30
1.4.2	Tools to Catch Coding Errors	31
1.5	Interview with Joshua Harlow	32
2	Modules, Libraries and Frameworks	36
2.1	The Import System	36
2.1.1	The sys Module	38
2.1.2	Import Paths	39
2.1.3	Custom Importers	40
2.1.4	Meta Path Finders	41
2.2	Useful Standard Libraries	43

CONTENTS

2.3	External Libraries	46
2.3.1	The External Libraries Safety Checklist	47
2.3.2	Protecting Your Code with an API Wrapper	48
2.4	Package Installation: Getting More From pip	49
2.5	Using and Choosing Frameworks	52
2.6	Doug Hellmann on Python Libraries	53
3	Documentation and Good API Practice	63
3.1	Documenting with Sphinx	64
3.2	Getting started with Sphinx and reST	66
3.3	Sphinx Modules	68
3.4	Automating the Table of Contents with autosummary	69
3.5	Automating Testing with doctest	70
3.6	Writing a Sphinx Extension	72
3.7	Managing Changes to Your API	74
3.7.1	Numbering API Versions	75
3.7.2	Documenting Your API Changes	76
3.7.3	Marking Deprecated Functions with the warnings Module	78
3.8	Summary	81
3.9	Interview with Christophe de Vienne	82
4	Handling Timestamps and Time Zones	87
4.1	The Problem of Missing Time Zones	88
4.2	Building Default datetime Objects	89

CONTENTS

4.3	Time Zone-Aware Timestamps with <code>dateutil</code>	90
4.4	Serializing datetime Objects	93
4.5	Solving Ambiguous Times	95
4.6	Summary	96
5	Distributing Your Software	98
5.1	A Bit of History	98
5.2	Packaging with <code>setup.cfg</code>	102
5.3	The <i>Wheel</i> Format Distribution Standard	104
5.4	Sharing Your Work with the World	107
5.5	Entry Points	112
5.5.1	Visualising Entry Points	113
5.5.2	Using Console Scripts	115
5.5.3	Using Plugins and Drivers	118
5.5.4	Summary	122
5.6	Interview with Nick Coghlan	123
6	Unit Testing	126
6.1	The Basics of Testing	126
6.1.1	Some Simple Tests	127
6.1.2	Skipping Tests	130
6.1.3	Running Particular Tests	132
6.1.4	Running Test in Parallel	134
6.1.5	Creating Objects Used in Tests with <code>fixtures</code>	135

CONTENTS

6.1.6	Running Test Scenarios	138
6.1.7	Controlled Test Using Mocking	139
6.1.8	Revealing Untested Code with coverage	145
6.2	Virtual Environments	149
6.2.1	Setting Up a Virtual Environment	150
6.3	Using virtualenv with tox	152
6.3.1	Re-creating an Environment	154
6.3.2	Using Different Python Versions	155
6.3.3	Integrating Other Tests	156
6.4	Testing Policy	157
6.5	Robert Collins on Testing	159
7	Methods and Decorators	163
7.1	Decorators and When to Use Them	163
7.1.1	Creating Decorators	164
7.1.2	Writing Decorators	165
7.1.3	Stacking Decorators	167
7.1.4	Writing Class Decorators	168
7.1.5	Retrieving Original Attributes with the <code>update_wrapper</code> Decorator	170
7.1.6	<code>wraps</code> : A Decorator for Decorators	172
7.1.7	Extracting Relevant Information with <code>inspect</code>	173
7.2	How Methods Work in Python	174
7.3	Static methods	176

CONTENTS

7.4	Class Methods	178
7.5	Abstract Methods	179
7.6	Mixing Static, Class, and Abstract Methods	181
7.7	Putting Implementations in Abstract Methods	184
7.8	The Truth About <code>super</code>	185
7.9	Summary	189
8	Functional Programming	190
8.1	Creating Pure Functions	191
8.2	Generators	192
8.2.1	Creating a Generator	193
8.2.2	Returning and Passing Values with <code>yield</code>	195
8.2.3	Inspecting Generators	197
8.3	List comprehensions	199
8.4	Functional Functions Functioning	200
8.4.1	Applying Functions to Items with <code>map</code>	201
8.4.2	Filtering Lists with <code>filter</code>	201
8.4.3	Getting Indexes with <code>enumerate</code>	202
8.4.4	Sorting a List with <code>sorted</code>	202
8.4.5	Finding Items That Satisfy Conditions with <code>any</code> and <code>all</code>	203
8.4.6	Combining Lists with <code>zip</code>	204
8.4.7	A Common Problem Solved	205
8.4.7.1	Finding the Item with Simple Code	205
8.4.7.2	Finding the Item Using <code>first</code>	207

CONTENTS

8.4.7.3	Using Lambda with <code>functools</code>	207
8.5	Using <code>itertools</code> Functions	210
8.6	Summary	212
9	The Abstract Syntax Tree, Hy, And Lisp-Like Attributes	213
9.1	Looking at the AST	214
9.2	Writing a Program Using the AST	216
9.2.1	The AST Objects	218
9.2.2	Walking Through an AST	219
9.3	Extending <code>flake8</code> with AST Checks	221
9.3.1	Writing the Class	222
9.3.2	Ignoring Irrelevant Code	223
9.3.3	Checking for the Correct Decorator	224
9.3.4	Looking for <code>self</code>	225
9.4	A Quick Introduction to Hy	228
9.5	Summary	231
9.6	Interview with Paul Tagliamonte	231
10	Performances and Optimizations	236
10.1	Data Structures	237
10.2	Understanding Behavior Through Profiling	240
10.2.1	<code>cProfile</code>	240
10.2.2	Disassembling with the <code>dis</code> Module	243
10.3	Defining Functions Efficiently	246

CONTENTS

10.4	Ordered Lists and Bisect	248
10.5	Namedtuple and Slots	252
10.6	Memoization	259
10.7	Faster Python with PyPy	261
10.8	Achieving Zero Copy with the Buffer Protocol	263
10.9	Summary	270
10.10	Victor Stinner on Optimization	270
11	Scaling and Architecture	274
11.1	Multithreading in Python and Its Limitations	274
11.2	Multiprocessing vs. Multithreading	277
11.3	Event-Driven Architecture	279
11.4	Other Options and <code>asyncio</code>	282
11.5	Service-Oriented Architecture	284
11.6	Interprocess Communication with ZeroMQ	285
11.7	Summary	288
12	Managing Relational Databases	289
12.1	RDBMS, ORMs, and When to Use Them	289
12.2	Database Backends	293
12.3	Streaming data with Flask and PostgreSQL	294
12.3.1	Writing the Data-Streaming Application	294
12.3.2	Building the Application	298
12.4	Dimitri Fontaine on Databases	301

13 Write Less, Code More	310
13.1 Using Six for Python 2 and 3 Support	310
13.1.1 Strings and Unicode	312
13.1.2 Handling Python Modules Moves	312
13.1.3 The modernize Module	313
13.2 Using Python Like Lisp to Make a Single Dispatcher	313
13.2.1 Creating Generic Methods in Lisp	314
13.2.2 Generic Methods in Python	316
13.3 Context Managers	319
13.4 Less Boilerplate with attr	323
Index	327

List of Figures

- 1.1 Python release timeline 22
- 1.2 Standard package directory 24
- 3.1 Explanation of some deprecated functions 78
- 6.1 Coverage of `ceilometer.publisher` 148
- 10.1 KCacheGrind example 242
- 10.2 Using slice on *memoryview* objects 266

List of Examples

1.1	<i>A pep8 run</i>	29
1.2	Running pep8 with <code>--ignore</code>	30
2.1	The Zen of Python	36
2.2	<i>Hy</i> module importer	41
2.3	<i>Hy</i> module loader	41
3.1	Building a basic Sphinx HTML document	67
3.2	Indicating the modules for autodoc to document	68
3.3	Code from <code>sphinxcontrib.pecanwsme.rest.setup</code>	73
3.4	An example of API change documentation for a car object	76
3.5	A documented change to the car object API using the warning module	79
3.6	Running python <code>-W error</code> and getting a deprecation error	80
3.7	An API change automated with <code>debtcollector</code>	81
4.1	Getting the time of the day	89
4.2	Constructing a datetime object from a date	90
4.3	Using <code>dateutil</code> objects as <code>tzinfo</code> classes	91
4.4	Obtaining your local time zone	92
4.5	Serializing a datetime object	93
4.6	Using the <code>iso8601</code> module to parse a ISO 8601 formatted timestamp	94
4.7	A confusing timestamp, occuring during the daylight saving time crossover	95
4.8	Disambiguating the ambiguous timestamp	96
5.1	<code>setup.py</code> using <i>distutils</i>	99
5.2	<code>setup.py</code> using <i>setuptools</i>	99
5.3	A <code>setup.cfg</code> file	102

LIST OF EXAMPLES

5.4	<code>setup.py</code> using <i>pbr</i>	103
5.5	Using <code>setup.py</code> <i>sdist</i>	108
5.6	Getting list of entry points groups	114
5.7	Showing details of an entry point group	114
5.8	<code>foobar/client.py</code>	116
5.9	<code>foobar/server.py</code>	116
5.10	A console script generated by <i>setuptools</i>	117
5.11	Running <i>pytimed</i>	120
5.12	Using <i>stevedore</i> to run a single extension from an entry point	122
6.1	A really simple test in <code>test_true.py</code>	127
6.2	A failing test in <code>test_true.py</code>	128
6.3	Skipping tests	131
6.4	Filtering tests run by name	132
6.5	Teardown functionality	136
6.6	Running a test using parameterized fixtures	138
6.7	Accessing the <code>mock.Mock</code> attribute	140
6.8	Creating methods on a <code>mock.Mock</code> object	140
6.9	Creating methods on a <code>mock.Mock</code> object with side effects	141
6.10	Checking method calls	141
6.11	Using <code>mock.patch</code>	142
6.12	Using <code>mock.patch</code> to test a set of behaviour	143
6.13	Using coverage with <i>pytest</i>	146
6.14	Automatic virtual environment creation	151
6.15	A <code>.travis.yml</code> example file	158
6.16	A <code>.travis.yml</code> example file with <i>tox-travis</i>	158
7.1	A decorator to organize functions in a dictionary	165
7.2	Adding a decorator to the factored code	166
7.3	Using more than one decorator with a single function	167
7.4	A decorated function loses its docstring and name attributes.	170
7.5	Source code of <code>functools.update_wrapper</code>	171

LIST OF EXAMPLES

7.6	Using <code>functools.wraps</code>	172
7.7	Retrieving function arguments using <code>inspect</code>	173
7.8	Calling bound <code>get_size</code>	176
7.9	Creating a static method as part of a class	177
7.10	A class method is bound to its class	178
7.11	Calling an abstract method	180
7.12	Implementing an abstract method using <code>abc</code>	181
7.13	Using a subclass to extend the signature of the abstract method of its parent	182
7.14	Using class method decorator with an abstract method	183
7.15	Using an implementation in an abstract method	184
7.16	The <code>super</code> function is a constructor that instantiates a <code>super object</code>	186
8.1	A non-pure function	191
8.2	A pure function	191
8.3	Creating a generator	193
8.4	<code>yield</code> returning a value	195
8.5	Checking whether a function is a generator	197
8.6	Source code of <code>inspect.isgeneratorfunction</code>	197
8.7	<i>map</i> usage in Python 3	201
8.8	<i>filter</i> usage in Python 3	201
8.9	Equivalent of <code>filter</code> using list comprehension	202
8.10	Returning a default value when the condition is not met	206
8.11	Finding the first item in a list that satisfies a condition	207
8.12	Using key without <code>lambda</code>	208
8.13	Using the <code>operator</code> module with <code>itertools.groupby</code>	211
9.1	Parsing Python code to AST	214
9.2	The AST of the <code>assign</code> command in Python	214
9.3	Compiling a Python AST	215
9.4	Hello world using Python AST	217
9.5	Walking a tree with <code>NodeTransformer</code> to alter a node	219

LIST OF EXAMPLES

9.6	Omitting and including <code>@staticmethod</code>	221
9.7	Registering new extensions for <code>flake8</code>	222
9.8	Template for <code>flake8</code> AST check	223
9.9	Ignoring statements that are not class definitions	223
9.10	Ignoring statements that are not function definitions	224
9.11	Checking for the static decorator	224
9.12	Checking the method for argumen	225
9.13	Checking the method for the self argument	226
9.14	Interacting with the Hy interpreter	229
9.15	Mapping a function definition from Python	229
9.16	Defining a class with <code>defclass</code>	229
9.17	Importing regular Python modules	230
9.18	Using <code>cond</code> instead of <code>if/elif/else</code>	230
10.1	Adding an entry in a dictionary of sets	238
10.2	Using the <code>cProfile</code> module	241
10.3	Using <i>KCacheGrind</i> to visualize Python profiling data	242
10.4	Disassembling Python code	243
10.5	Disassembling functions that concatenate strings	244
10.6	Disassembling a function defined in a function	246
10.7	Disassembling a closure	247
10.8	Using <code>bisect</code> to find a needle in a haystack	248
10.9	Inserting an item in a sorted list	249
10.10A	SortedList implementation	249
10.11	How attributes are stored internally in a Python object	252
10.12	Memory profiling with objects	253
10.13	Extract of <i>CPython</i> <code>Objects/typeobject.c</code> code	254
10.14A	class declaration using <code>__slots__</code>	255
10.15	Memory usage of objects using <code>__slots__</code>	256
10.16	Declaring a class using <code>namedtuple</code>	257
10.17	Memory usage of a class built from <code>collections.namedtuple</code>	258

LIST OF EXAMPLES

10.18	A basic memoization	259
10.19	Using <code>functools.lru_cache</code>	260
10.20	Writing a file with random data	263
10.21	Running memory profiler	264
10.22	Using <code>memoryview</code>	265
10.23	Writing random data to a file, using less memory	266
10.24	Running memory profiler with the optimized version	266
10.25	Reading data from a socket	267
10.26	Reading data from a socket using <code>memoryview</code>	268
10.27	Reading data from a file into a <code>bytearray</code>	269
10.28	Reading data from a file into a <code>bytearray</code> at a specified location	269
11.1	Result of <i>time python worker.py</i>	278
11.2	Using multiprocessing for concurrent activity	278
11.3	Result of <i>time python worker.py</i>	279
11.4	Event-driven program that listens for and processes connections	281
11.5	Retrieving web pages concurrently with <code>aiohttp</code>	283
12.1	Detecting and excluding duplicate messages with an ORM	291
12.2	SQL table schema for storing messages	294
12.3	The <i>notify_on_insert</i> function	295
12.4	The trigger for <i>notify_on_insert</i>	296
12.5	Receiving notifications in Python	296
12.6	Flask streamer application	298
13.1	Using <code>six.moves</code> to use <code>ConfigParser()</code> with Python 2 and Python 3	313
13.2	Defining generic methods in Lisp, independent of classes	315
13.3	Calling a method with an unavailable signature	316
13.4	A Python implementation of our drum kit	317
13.5	Simple implementation of a context object	319
13.6	Simplest usage of <code>contextlib.contextmanager</code>	320
13.7	Defining a context manager yielding a value	321
13.8	Handling exceptions in a context manager	322

LIST OF EXAMPLES

13.9 Opening two files at the same time to copy content	323
13.10Opening two files at the same time with one with statement	323
13.11Common class initialization boilerplate	323
13.12Declaring a class with attr	324
13.13Using attr.ib() with its converter argument	325
13.14Using frozen=True with attr	326

Introduction

If you’re reading this, the odds are good you’ve been working with **Python** for some time already. Maybe you learned it using some tutorials, delved into some existing programs, or started from scratch. Whatever the case, you’ve hacked your way into learning it. That’s exactly how I got familiar with Python up until I started working on big open source projects 10 years ago.

It is easy to think that you know and understand Python once you’ve written your first program. The language is that simple to grasp. However, it takes years to master it and to develop a deep comprehension of its advantages and shortcomings.

When I started Python, I built my own Python libraries and applications on a “garage project” scale. Things changed once I started working with hundreds of developers on software that thousands of users rely on. For example, the OpenStack platform — a project I contribute to — represents over 9 million lines of Python code, which collectively needs to be concise, efficient, and scalable to the needs of whatever cloud computing application its users require. When you have a project of this size, things like testing and documentation absolutely require automation, or else they won’t get done at all.

I thought I knew a lot about Python before working on projects of this scale — a scale I could hardly imagine when I started out — but I’ve learned a lot more. I’ve also had the opportunity to meet some of the best Python hackers in the industry and learn from them. They’ve taught me everything from general architecture and design principles to various helpful tips and tricks. Through this book, I hope to share the most important things I’ve learned so that you can build better Python programs — and build them more efficiently, too!

The first version of this book, *The Hacker's Guide to Python*, came out in 2014. Now *Serious Python* is the fourth edition, with updated and entirely new contents. I hope you enjoy it!

Who Should Read This Book and Why

This book is intended for Python coders and developers who want to take their Python skills to the next level.

In it, you'll find methods and advice that will help you get the most out of Python and build future-proof programs. If you're already working on a project, you'll be able to apply the techniques discussed right away to improve your current code. If you're starting your first project, you'll be able to create a blueprint with the best practice.

I'll introduce you to some Python internals to give you a better understanding of how to write efficient code. You will gain a greater insight into the inner workings of the language that will help you understand problems or inefficiencies.

The book also provides applicable battle-tested solutions to problems such as testing, porting, and scaling Python code, applications, and libraries. This will help you avoid making the mistakes that others have made and discover strategies that will help you maintain your software in the long run.

About this Book

This book is not necessarily designed to be read from front to back. You should feel free to skip to sections that interest you or are relevant to your work. Throughout the book, you'll find a wide range of advice and practical tips. Here's a quick breakdown of what each chapter contains.

Chapter 1 will provide guidance about what to consider before you undertake a project, with advice on structuring your project, numbering versions, setting up automated error checking, and more.

Chapter 2 will look at Python modules, libraries, and frameworks and talk a little about how they work under the hood. You'll find guidance on using the `sys` module, getting more from the `pip` package manager, choosing the best framework for you, and using standard and external libraries. There's also an interview with Doug Hellmann.

Chapter 3 gives advice on documenting your projects and managing your APIs as your project evolves even after publication. You'll get specific guidance on using Sphinx to automate certain documentation tasks. Here you'll find an interview with Christophe de Vienne.

Chapter 4 covers the age-old issue of time zones and how best to handle them in your programs using `datetime` objects and `tzinfo` objects.

Chapter 5 will help you get your software to users with guidance on distribution. You'll learn about packaging, distributions standards, the `distutils` and `setuptools` libraries, and how to easily discover dynamic features in a package using entry points. Nick Coghlan is interviewed.

Chapter 6 advises you on unit testing with best-practice tips and specific tutorials on automating unit tests with `pytest`. You'll also look at using virtual environments to increase the isolation of your tests. The interview is with Robert Collins.

Chapter 7 digs into methods and decorators. This is a look at using Python for functional programming, with advice on how and when to use decorators and how to create decorators for decorators. We'll also dig into static, class, and abstract methods and how to mix the three for a more robust program.

Chapter 8 will show you more functional programming tricks you can implement in Python. This chapter discusses generators, list comprehensions, functional functions and common tools for implementing them, and the useful `functools` li-

brary.

Chapter 9 peeks under the hood of the language itself and discusses the abstract syntax tree (AST) that is the inner structure of Python. We'll also look at extending `flake8` to work with the AST to introduce more sophisticated automatic checks into your programs. The chapter concludes with an interview with Paul Tagliamonte.

Chapter 10 is a guide to optimizing performance by using appropriate data structures, defining functions efficiently, and applying dynamic performance analysis to identify bottlenecks in your code. We'll also touch on memoization and reducing waste in data copies. You'll find an interview with Victor Stinner.

Chapter 11 tackles the difficult subject of multithreading, including how and when to use multithreading as opposed to multiprocessing and whether to use event-oriented or service-oriented architecture to create scalable programs.

In **Chapter 12**, we'll tackle relational databases. We'll take a look at how they work and how to use PostgreSQL to effectively manage and stream data. Dimitri Fontaine is interviewed.

Finally, in **Chapter 13**, you'll find sound advice on a range of topics: making your code compatible with both Python 2 and 3, creating functional Lisp-like code, using context managers, and reducing repetition with the `attr` library.

CHAPTER 1



Starting Your Project

In this first chapter, we'll look at a few aspects of starting out a project and what you should think about before you begin, such as which Python version to use, how to structure your modules, how to effectively number software versions, and how to ensure best coding practices with automatic error checking.

1.1 Versions of Python

Before beginning a project, you'll need to decide what version(s) of Python it will support. This is not as simple a decision as it may seem.

It's no secret that Python supports several versions at the same time. Each minor version of the interpreter gets bug-fix support for 18 months and security support for 5 years. For example, Python 3.7, released on June 27, 2018, will be supported until Python 3.8 is released, around October 2019 (15 months later). Around December 2019, a last bug-fix release of Python 3.7 will occur, and everyone will be expected to switch to Python 3.8. Each new version of Python introduces new features and deprecates old ones. Figure [1.1](#) illustrates this timeline.

1.1. VERSIONS OF PYTHON




Figure 1.1: Python release timeline

On top of that, we should take into consideration the Python 2 versus Python 3 problem. People working with (very) old platforms may still require Python 2 support because Python 3 has not been made available on those platforms, but the rule of thumb is to forget Python 2 if you can.

Here is a quick way to figure out which version you need:

- Versions 2.6 and older are now obsolete, so I do not recommend you worry about supporting them at all. If you do intend to support these older versions for whatever reason, be warned that you'll have a hard time ensuring that your program supports Python 3.x as well. Having said that, you might still run into Python 2.6 on some older systems — if that's the case, sorry!
- Version 2.7 is and will remain the last version of Python 2.x. Every system is basically running or able to run Python 3 one way or the other nowadays, so unless you're doing archeology, you shouldn't need to worry about supporting Python 2.7 in new programs. Python 2.7 will cease to be supported after the year 2020, so the last thing you want to do is build a new software based on it.
- Versions 3.7 is the most recent version of the Python 3 branch as of this writing, and that's the one that you should target. Most recent operating systems ship at least 3.6, so in the case where you'd target those, you can make sure your application also work with 3.7.

Techniques for writing programs that support both Python 2.7 and 3.x will be discussed in [Section 13.1](#).

Finally, note that this book has been written with Python 3 in mind.

1.2 Laying Out Your Project

Starting a new project is always a bit of a puzzle. You can't be sure how your project will be structured, so you might not know how to organize your files. However, once you have a proper understanding of best practices, you'll understand which basic structure to start with. Here I'll give some tips on dos and don'ts for laying out your project.

1.2.1 What to Do

First, consider your project structure, which should be fairly simple. Use packages and hierarchy wisely: a deep hierarchy can be a nightmare to navigate, while a flat hierarchy tends to become bloated.

Then, avoid making the common mistake of storing unit tests outside the package directory. These tests should definitely be included in a subpackage of your software so that they aren't automatically installed as a `tests` top-level module by `setuptools` (or some other packaging library) by accident. By placing them in a subpackage, you ensure they can be installed and eventually used by other packages so users can build their own unit tests.

[Figure 1.2](#) illustrates what a standard file hierarchy should look like.




Figure 1.2: Standard package directory

The standard name for a Python installation script is `setup.py`. It comes with its companion `setup.cfg`, which should contain the installation script configuration details. When run, `setup.py` will install your package using the Python distribution utilities.

You can also provide important information to users in `README.rst` (or `README.txt`, or whatever filename suits your fancy). Finally, the `docs` directory should contain

the package’s documentation in *reStructuredText* format, which will be consumed by *Sphinx* (see Section 3.2).

Packages will often have to provide extra data for the software to use, such as images, shell scripts, and so forth. Unfortunately, there’s no universally accepted standard for where these files should be stored, so you should just put them wherever makes the most sense for your project depending on their functions. For example, web application templates could go in a `templates` directory in your package root directory.

The following top-level directories also frequently appear:

- `etc` for sample configuration files
- `tools` for shell scripts or related tools
- `bin` for binary scripts you’ve written that will be installed by `setup.py`

1.2.2 What Not to Do

There is a particular design issue that I often encounter in project structures that have not been fully thought out: some developers will create files or modules based on the type of code they will store. For example, they might create `functions.py` or `exceptions.py` files. This is a terrible approach and doesn’t help any developer when navigating the code. When reading a code base, the developer expects a functional area of a program to be confined in a particular file. The code organization doesn’t benefit from this approach, which forces readers to jump between files for no good reason.

Organize your code based on features, not on types.

It is also a bad idea to create a module directory that contains only an `__init__.py` file, because it’s unnecessary nesting. For example, you shouldn’t create a directory named `hooks` with single file named `hooks/__init__.py` in it, where `hooks.py` would

have been enough. If you create a directory, it should contain several other Python files that belong to the category the directory represents. Building a deep hierarchy unnecessarily is confusing.

You should also be very careful about the code that you put in the `__init__.py` file: this file will be called and executed the first time that a module contained in the directory is loaded. Placing the wrong things in your `__init__.py` can have unwanted side effects. In fact, `__init__.py` files should be empty most of the time, unless you know what you're doing. Don't try to remove `__init__.py` files altogether though, or you won't be able to import your Python module at all: Python requires an `__init__.py` file to be present for the directory to be considered a submodule.

1.3 Version Numbering

Software versions need to be stamped so users know which is the more recent version. For every project, users must be able to organize the timeline of the evolving code.

There is an infinite number of ways to organize your version numbers. However, [PEP 440](#) introduces a version format that every Python package, and ideally every application, should follow so that other programs and packages can easily and reliably identify which versions of your package they require.

PEP 440 defines the following regular expression format for version numbering:

```
1 N[.N]+[{a|b|c|rc}N][.postN][.devN]
```

This allows for standard numbering such as 1.2 or 1.2.3. There are a few further details to note:

- Version 1.2 is equivalent to 1.2.0, 1.3.4 is equivalent to 1.3.4.0, and so forth.
- Versions matching `N[.N]+` are considered *final* releases.

1.3. VERSION NUMBERING

- Date-based versions such as 2013.06.22 are considered invalid. Automated tools designed to detect PEP 440-format version numbers will (or should) raise an error if they detect a version number greater than or equal to 1980.

Final components can also use the following format:

- $N[.N]+aN$ (for example, 1.2a1) denotes an **alpha** release; this is a version that might be unstable and missing features.
- $N[.N]+bN$ (for example, 2.3.1b2) denotes a **beta** release, a version that might be feature-complete but still buggy.
- $N[.N]+cN$ or $N[.N]+rcN$ (for example, 0.4rc1) denotes a (release) candidate. This is a version that might be released as the final product unless significant bugs emerge. While the rc and c suffixes have the same meaning, if both are used, rc releases are considered to be newer than c releases.

The following suffixes can also be used:

- .postN (for example, 1.4.post2) indicates a post release. Post releases are typically used to address minor errors in the publication process, such as mistakes in release notes. You shouldn't use .postN when releasing a bug-fix version; instead but should increment the minor version number.
- .devN (for example, 2.3.4.dev3) indicates a developmental release. It indicates a prerelease of the version that it qualifies: for example, 2.3.4.dev3 indicates the third developmental version of the 2.3.4 release, prior to any alpha, beta, candidate or final release. This suffix is discouraged because it is harder for humans to parse.

This scheme should be sufficient for most common use cases.

Note

You might have heard of **Semantic Versioning**, which provides its own guidelines for version numbering. This specification partially overlaps with PEP 440, but unfortunately, they're not entirely compatible. For example, Semantic Versioning's recommendation for prerelease versioning uses a scheme such as `1.0.0-alpha+001` that is not compliant with PEP 440.

Many *distributed version control system (DVCS)* platforms, such as Git and Mercurial, are able to generate version numbers using an identifying hash (for Git, refer to `git describe`). Unfortunately, this system isn't compatible with the scheme defined by PEP 440: for one thing, identifying hashes aren't orderable.

1.4 Coding Style & Automated Checks

Coding style is a touchy subject, but one we should talk about before we dive further into Python. Unlike many programming languages, Python uses indentation to define blocks. While this offers a simple solution to the age-old question “Where should I put my braces?” it introduces a new question: “How should I indent?”

That was one of the first questions raised in the community, so the Python folks, in their vast wisdom, came up with the **PEP 8: Style Guide for Python Code**.

This document defines the standard style for writing Python code. The list of guidelines boils down to:

- Use four spaces per indentation level.
- Limit all lines to a maximum of 79 characters.
- Separate top-level function and class definitions with two blank lines.

1.4. CODING STYLE & AUTOMATED CHECKS

- Encode files using ASCII or UTF-8.
- One module import per `import` statement and per line. Place import statements at the top of the file, after comments and docstrings, grouped first by standard, then by third party, and finally by local library imports.
- Do not use extraneous whitespaces between parentheses, brackets, or braces, or before commas.
- Write class names in camel case (e.g, `CamelCase`), suffix exceptions with `Error` (if applicable), and name functions in lowercase with words and underscores (e.g, `separated_by_underscores`). Use a leading underscore for `_private` attributes or methods.

These guidelines really aren't hard to follow, and they make a lot of sense. Most Python programmers have no trouble sticking to them as they write code.

However, *errare humanum est* and it's still a pain to look through your code to make sure it fits the PEP 8 guidelines. Luckily, there's a [pep8](#) tool that can automatically check any Python file you send its way. Install `pep8` with `pip`, and then you can use it on a file like so:

Example 1.1 A *pep8* run

```
1 $ pep8 hello.py
2 hello.py:4:1: E302 expected 2 blank lines, found 1
3 $ echo $?
4 1
```

Here I use `pep8` on my file `hello.py`, and the output indicates which lines and columns do not conform to PEP 8 and reports each issue with a code — here it's line 4 and column 1. Violations of *MUST* statements in the specification are reported as errors, and their error codes start with an `E`. Minor issues are reported as warnings, and their error codes start with a `W`. The three-digit code following that first letter indicates the exact kind of error or warning.

The hundreds digit tells you the general category of an error code: for example, errors starting with E2 indicate issues with whitespace, errors starting with E3 indicate issues with blank lines, and warnings starting with W6 indicate deprecated features being used. These codes are all listed in the [pep8 documentation](#).

1.4.1 Tools to Catch Style Errors

The community still debates whether validating against PEP 8 code, which is not part of the standard library, is good practice. My advice is to consider running a PEP 8 validation tool against your source code on a regular basis. You can do this easily by integrating it into your continuous integration system. While this approach may seem a bit extreme, it's a good way to ensure that you continue to respect the PEP 8 guidelines in the long term. We'll discuss in “Using virtualenv with tox” in Section 6.3 how you can integrate pep8 with tox to automate these checks.

Most open source projects enforce PEP 8 conformance through automatic checks. Using these automatic checks from the very beginning of the project might frustrate newcomers, but it also ensures that the codebase always looks the same in every part of the project. This is very important for a project of any size where there are multiple developers with differing opinions on, for example, whitespace ordering. You know what I mean.

It's also possible to set your code to ignore certain kinds of errors and warnings by using the `--ignore` option, like so:

Example 1.2 Running pep8 with `--ignore`

```
1 $ pep8 --ignore=E3 hello.py
2 $ echo $?
3 0
```

This will ignore any code E3 errors inside my `hello.py` file. The `--ignore` option allows you to effectively ignore parts of the PEP 8 specification that you don't want

to follow. If you're running `pep8` on an existing codebase, it also allows you to ignore certain kinds of problems so you can focus on fixing issues one category at a time.



Note

If you write C code for Python (e.g. modules), the [PEP 7](#) standard describes the coding style that you should follow.

1.4.2 Tools to Catch Coding Errors

Python also has tools that check for actual coding errors rather than style errors. Here are some notable examples:

- [pyflakes](#): extendable via plugins.
- [pylint](#): Checks PEP 8 conformance while performing code error checks by default; can be extended via plugin.

These tools all make use of static analysis — that is, they parse the code and analyze it rather than running it outright.

If you choose to use `Pyflakes`, note that it doesn't check PEP 8 conformance on its own, so you'd need the second `pep8` tool to cover both.

To simplify things, Python has a project named [flake8](#) that combines `pyflakes` and `pep8` into a single command. It also adds some new fancy features: for example, it can skip checks on lines containing `#noqa` and is extensible via plugins.

There are a large number of plugins available for `flake8` that you can use out the box. For example, installing `flake8-import-order` (with `pip install flake8-import-order`) will extend `flake8` so that it also checks whether your `import` statements are sorted alphabetically in your source code. Yes, some projects want that.

In most open source projects, `flake8` is heavily used for code style verification. Some large open source projects have even written their own plugins for `flake8`, adding checks for errors such as odd usage of `except`, Python 2/3 portability issues, import style, dangerous string formatting, possible localization issues, and more.

If you're starting a new project, I strongly recommend that you use one of these tools for automatic checking of your code quality and style. If you already have a codebase that didn't implement automatic code checking, a good approach is to run your tool of choice with most of the warnings disabled and fix issues one category at a time.

Though none of these tools may be a perfect fit for your project or your preferences, `flake8` is a good way to improve the quality of your code and make it more durable.



Tip

Many text editors, including the famous **GNU Emacs** and **vim**, have plugins available (such as *Flycheck* that can run tools such as `pep8` or `flake8` directly in your code buffer, interactively highlighting any part of your code that isn't PEP 8 compliant. This is a handy way to fix most style errors as you write your code.

We'll talk about extending this toolset in Section [9.3](#) with our own plugin to verify correct method declaration.

1.5 Interview with Joshua Harlow

Joshua Harlow is a Python developer. He was one of the technical leads on the OpenStack team at Yahoo! between 2012 and 2016 and now works at GoDaddy. Josh is the author of several Python libraries such as `Taskflow`, `automaton`, and `Zake`.

“

**What got you into using Python?**

I started programming in Python 2.3 or 2.4 back in about 2004 during an internship at the IBM near Poughkeepsie, New York (most of my relatives and family are from upstate NY, shout out to them!). I forget exactly what I was doing there, but it involved wxPython and some Python code that they were working on to automate some system.

After that internship I returned to school, went on to graduate school at the Rochester Institute of Technology (RIT), and ended up working at Yahoo!.

I eventually ended up in the CTO team, where I and a few others were tasked with figuring out which open source cloud platform to use. We landed on OpenStack, which is written almost entirely in Python.

What do you love and hate about the Python language?

Some of the things I love (not a comprehensive listing):

- Its simplicity — Python is really easy for beginners to engage with and for experienced developers to stay engaged with.
- Style checking—reading code you wrote later on is a big part of developing software and having consistency that can be enforced by tools such as `flake8`, `pep8`, and `PyLint` really helps.

- The ability to pick and choose programming styles and mix them up as you see fit.

Some of the things I dislike (not a comprehensive listing):

- The somewhat painful Python 2 to 3 transition (version 3.6 has paved over most of the issues here).
- Lambdas are too simplistic and should be made more powerful.
- The lack of a decent package installer — I feel pip needs some work, like developing a real dependency resolver.
- The global interpreter lock (GIL) and the need for it. It makes me sad.
- The lack of native support for multithreading — currently you need the addition of an explicit asyncio model.
- The fracturing of the Python community; this is mainly around the split between CPython and PyPy (and other variants).

You work on `debtcollector`, a Python module for managing deprecation warnings. How is the process of starting a new library?

The simplicity mentioned above makes it really easy to get a new library going and to publish it so others can use it. Since that code came out of one of the other libraries that I work on (`taskflow`) it was relatively easy to transplant and extend that code without having to worry about the API being badly designed. I am very glad others (inside the OpenStack community or outside of it) have found a need/use for it, and I hope that library grows to accommodate more styles of deprecation patterns that other libraries (and applications?) find useful.

What is Python missing, in your opinion?

Python could perform better under just-in-time (JIT) compilation. Most newer languages being created (such as Rust, Node.js using the Chrome V8 JavaScript engine, and others) have many of Python's capabilities but are also JIT compiled. It would really be great if the default CPython could also be JIT compiled so that Python could compete with these newer languages on performance.

Python also really needs a strong set of concurrency patterns; not just the low level asyncio and threading styles of patterns, but higher-level concepts that help make applications that work performantly at larger scale. The Python library goesl does port over some of the concepts from Go, which does provide a built-in concurrency model. I believe these higher-level patterns need to be available as first-class patterns that are built in to the standard library and maintained so that developers can use them where they see fit. Without these, I don't see how Python can compete with other languages that do provide them.

Until next time, keep coding and be happy!

”

CHAPTER 2



Modules, Libraries and Frameworks

Modules are an essential part of what makes Python extensible. Without them, Python would just be a language built around a monolithic interpreter; it wouldn't have flourish within a giant ecosystem that allows developers to build applications quickly and simply by combining extensions. In this chapter, I'll introduce you to some of the features that make Python modules great, from the built-in modules you need to know to externally managed frameworks.

2.1 The Import System

To use modules and libraries in your programs, you have to import them using the `import` keyword. As an example, Example 2.1 imports the all-important Zen of Python guidelines:.

Example 2.1 The Zen of Python

2.1. THE IMPORT SYSTEM

```
1 >>> import this
2 The Zen of Python, by Tim Peters
3
4 Beautiful is better than ugly.
5 Explicit is better than implicit.
6 Simple is better than complex.
7 Complex is better than complicated.
8 Flat is better than nested.
9 Sparse is better than dense.
10 Readability counts.
11 Special cases aren't special enough to break the rules.
12 Although practicality beats purity.
13 Errors should never pass silently.
14 Unless explicitly silenced.
15 In the face of ambiguity, refuse the temptation to guess.
16 There should be one-- and preferably only one --obvious way to do it.
17 Although that way may not be obvious at first unless you're Dutch.
18 Now is better than never.
19 Although never is often better than *right* now.
20 If the implementation is hard to explain, it's a bad idea.
21 If the implementation is easy to explain, it may be a good idea.
22 Namespaces are one honking great idea -- let's do more of those!
```

The import system is quite complex, and I'm assuming you already know the basics, so here I'll show you some of the internals of this system, including the workings of the sys module, how to change or add import paths, and using custom importers.

First, you need to know that the import keyword is actually a wrapper around a function named `__import__`. Here is a familiar way of importing a module:

```
1 >>> import itertools
2 >>> itertools
3 <module 'itertools' from '/usr/.../>
```

This is precisely equivalent to this method:

```
1 >>> itertools = __import__("itertools")
```

2.1. THE IMPORT SYSTEM

```
2 >>> itertools
3 <module 'itertools' from '/usr/.../>
```

You can also imitate the `as` keyword of `import`, as those two equivalent ways of importing show:

```
1 >>> import itertools as it
2 >>> it
3 <module 'itertools' from '/usr/.../>
```

And here's the second example:

```
1 >>> it = __import__("itertools")
2 >>> it
3 <module 'itertools' from '/usr/.../>
```

The `__import__` function is extremely useful to know, as in some (corner) cases, you might want to import a module whose name is unknown beforehand, like so:

```
1 >>> random = __import__("RANDOM".lower())
2 >>> random
3 <module 'random' from '/usr/.../>
```

Don't forget that module, once imported, are essentially object whose attributes (classes, functions, variables, etc) are objects.

2.1.1 The `sys` Module

The `sys` module provides access to variables and functions related to Python itself and the operating system it is running on. This module also contains a lot of information about Python's import system.

First of all, you can retrieve the list of modules currently imported using the `sys.modules` variable. The `sys.modules` variable is a dictionary whose key is the module name you want to inspect and whose returned value is the module object. For example, once the `os` module is imported, we can retrieve it by entering:

2.1. THE IMPORT SYSTEM

```
1 >>> import sys
2 >>> import os
3 >>> sys.modules['os']
4 <module 'os' from '/usr/lib/python2.7/os.pyc'>
```

The `sys.modules` variable is a standard Python dictionary that contains all loaded modules. That means that calling `sys.modules.keys()`, for example, will return the complete list of the names of loaded modules.

You can also retrieve the list of modules that are built in by using the `sys.builtin_module_names` variable. The built-in modules compiled to your interpreter can vary depending on what compilation options were passed to the Python build system.

2.1.2 Import Paths

When importing modules, Python relies on a list of paths to know where to look for the module. This list is stored in the `sys.path` variable. To check which paths your interpreter will search for modules, just enter `sys.path`.

You can change this list, adding or removing paths as necessary, or even modify the `PYTHONPATH` environment variable to add paths, without writing Python code at all. Adding paths to the `sys.path` variable can be useful if you want to install Python modules to non-standard locations, such as a test environment. In normal operations, however, it should not be necessary to change the path variable. The following approaches are almost equivalent – *almost* because the path will not be placed at the same level in the list; this difference may not matter, depending on your use case:

```
1 >>> import sys
2 >>> sys.path.append('/foo/bar')
```

This would be (almost) the same as:

```
1 $ PYTHONPATH=/foo/bar python
2 >>> import sys
3 >>> '/foo/bar' in sys.path
4 True
```

It's important to note that the list will be iterated over to find the requested module, so the order of the paths in `sys.path` is important. It's useful to put the path most likely to contain the modules you are importing early in the list to speed up search time. Doing so also ensures that if two modules with the same name are available, the first match will be picked.

This last property is especially important because one common mistake is to shadow Python built-in modules with your own. Your current directory is searched before the Python Standard Library directory. That means that if you decide to name one of your scripts `random.py` and then try using `import random`, the file from your current directory will be imported and not the Python module.

2.1.3 Custom Importers

You can also extend the import mechanism using custom importers. This is the technique that the Lisp-Python dialect Hy uses to teach Python how to import files other than standard `.py` or `.pyc` files. Hy is a Lisp implementation on top of Python, discussed later in Section [9.4](#).

The import hook mechanism, as this technique is called, is defined by [PEP 302](#). It allows you to extend the standard import mechanism, which in turn can allow you to modify how Python imports modules and build your own system of import. For example, you could write an extension that imports modules from a database over the network, or that would do some sanity checking before importing any module. Python offers two different but related ways to broaden the import system: the meta path finders, for use with `sys.meta_path`, and the path entry finders for use with `sys.path_hooks`.

2.1.4 Meta Path Finders

The meta path finder is an object that will allow you to load custom objects as well as the regular and standard .py files that Python knows how to load. A meta path finder object must expose a `find_module(fullname, path=None)` method that returns a loader object. The loader object must also have a `load_module(fullname)` method responsible for loading the module from a source file. To illustrate, Example 2.2 shows how Hy uses a custom meta path finder to enable Python to import source files ending with .hy instead of .py.

Example 2.2 *Hy* module importer

```

1 class MetaImporter(object):
2     def find_on_path(self, fullname):
3         fls = ["%s/__init__.hy", "%s.hy"]
4         dirpath = "/".join(fullname.split("."))
5
6         for pth in sys.path:
7             pth = os.path.abspath(pth)
8             for fp in fls:
9                 composed_path = fp % ("%s/%s" % (pth, dirpath))
10                if os.path.exists(composed_path):
11                    return composed_path
12
13     def find_module(self, fullname, path=None):
14         path = self.find_on_path(fullname)
15         if path:
16             return MetaLoader(path)
17
18 sys.meta_path.append(MetaImporter())

```

Once Python has determined that the path is valid and that it points to a module, a `MetaLoader` object is returned as shown in Example 2.3.

Example 2.3 *Hy* module loader

2.1. THE IMPORT SYSTEM

```
1 class MetaLoader(object):
2     def __init__(self, path):
3         self.path = path
4
5     def is_package(self, fullname):
6         dirpath = "/" .join(fullname.split("."))
7         for pth in sys.path:
8             pth = os.path.abspath(pth)
9             composed_path = "%s/%s/__init__.hy" % (pth, dirpath)
10            if os.path.exists(composed_path):
11                return True
12        return False
13
14    def load_module(self, fullname):
15        if fullname in sys.modules:
16            return sys.modules[fullname]
17
18        if not self.path:
19            return
20
21        sys.modules[fullname] = None
22        mod = import_file_to_module(fullname,
23                                   self.path) ❶
24
25        ispkg = self.is_package(fullname)
26
27        mod.__file__ = self.path
28        mod.__loader__ = self
29        mod.__name__ = fullname
30
31        if ispkg:
32            mod.__path__ = []
33            mod.__package__ = fullname
34        else:
35            mod.__package__ = fullname.rpartition('.')[0]
```

```

37 sys.modules[fullname] = mod
38 return mod

```

- ❶ `import_file_to_module` reads a *Hy* source file, compiles it to Python code, and returns a Python module object.

This loader is pretty straightforward actually: once the `.hy` file is found, it's passed to this loader, which compiles the file if necessary, registers it, sets some attributes, and then returns it to the Python interpreter.

The `uprefix` module is another good example of this feature in action. Python 3.0 through 3.2 didn't support the `u` prefix for denoting Unicode strings that was featured in Python 2; the `uprefix` module ensures compatibility between Python versions 2 and 3 by removing the `u` prefix from strings before compilation.

2.2 Useful Standard Libraries

Python comes with a huge standard library packed with tools and features for almost any purpose you can think of. Newcomers to Python who are used to having to write their own functions for basic tasks are often shocked to find that the language itself ships with so much functionality built in and ready for use.

Whenever you're tempted to write your own function to handle a simple task, first stop and look through the standard library. In fact, My advice is to skim through the whole thing at least once before you begin working with Python so that next time you need a function, you have an idea of whether it already exists in the standard library. We'll talk about some of these modules in later sections, such as `functools` and `itertools`, in later chapters, but here are a few of the standard modules that you'll definitely find useful:

- **`atexit`** allows you to register functions to call when your program exits.

2.2. USEFUL STANDARD LIBRARIES

- **argparse** provides functions for parsing command line arguments.
- **bisect** provides bisection algorithms for sorting lists (see Section [10.4](#)).
- **calendar** provides a number of date-related functions.
- **codecs** provides functions for encoding and decoding data.
- **collections** provides a variety of useful data structures.
- **copy** provides functions for copying data.
- **csv** provides functions for reading and writing CSV files.
- **datetime** provides classes for handling dates and times.
- **fnmatch** provides functions for matching Unix-style filename patterns.
- **concurrent** provides asynchronous computation (native in Python 3, available for Python 2 via PyPI).
- **glob** provides functions for matching Unix-style path patterns.
- **io** provides functions for handling I/O streams. It also contains **StringIO** which allows you to treat strings as files.
- **json** provides functions for reading and writing data in JSON format.
- **logging** provides access to Python's own built-in logging functionality.
- **multiprocessing** allows you to run multiple subprocesses from your application, while providing an API that makes them look like threads.
- **operator** provides functions implementing the basic Python operators, which you can use instead of having to write your own lambda expressions (see Section [8.4](#)).
- **os** provides access to basic OS functions.

2.2. USEFUL STANDARD LIBRARIES

- **random** provides functions for generating pseudo-random numbers.
- **re** provides regular expression functionality.
- **sched** provides an event scheduler without using multi-threading.
- **select** provides access to the *select()* and *poll()* functions for creating event loops.
- **shutil** provides access to high-level file functions.
- **signal** provides functions for handling POSIX signals.
- **tempfile** provides functions for creating temporary files and directories.
- **threading** provides access to high-level threading functionality.
- **urllib** (and **urllib2** and **urlparse** in Python 2.x) provides functions for handling and parsing URLs.
- **uuid** allows you to generate UUIDs (Universally Unique Identifiers).

Use this list as a quick reference for what these useful libraries modules do. If you can memorize even part of this list, all the better. The less time you have to spend looking up library modules, the more time you can spend writing the code you actually need.

Most of the standard library is written in Python, so there's nothing stopping you from looking at the source code of the modules and functions. When in doubt, crack open the code and see what it does for yourself. Even if the documentation has everything you need to know, there's always a chance you could learn something useful.

2.3 External Libraries

Python’s “batteries included” philosophy is that, once you have Python installed, you should have everything you need to build whatever you want. This is to prevent the programming equivalent of unwrapping an awesome gift only to find out that whoever gave it to you forgot to buy batteries for it.

Unfortunately, there’s no way the people behind Python can predict *everything* you might want to make. And even if they could, most people wouldn’t want to deal with a multi-gigabyte download, especially if they just wanted to write a quick script for renaming files. So, even with all its extensive functionality, the Python Standard Library doesn’t cover everything. Luckily, members of the Python community have created external libraries.

The Python Standard Library is safe, well-charted territory: its modules are heavily documented, and enough people use it on a regular basis that you can feel assured that it won’t break messily when you give it a try – and in the unlikely event that it *does* break, you can be confident someone will fix it in short order. External libraries, on the other hand, are the parts of the map labeled “here there be dragons”: documentation may be sparse, functionality may be buggy, and updates may be sporadic or even nonexistent. Any serious project will likely need functionality that only external libraries can provide, but you need to be mindful of the risks involved in using them.

Here’s a tale of external library dangers from the trenches. OpenStack uses **SQLAlchemy** – a database toolkit for Python. If you’re familiar with SQL, you know that database schemas can change over time, so OpenStack also made use of **sqlalchemy-migrate** to handle schema migration needs. And it worked... until it didn’t. Bugs started piling up, and nothing was getting done about them. At this time, OpenStack was also interested in supporting Python 3, but there was no sign that sqlalchemy-migrate was moving toward Python 3 support. It was clear by that point that sqlalchemy-migrate was effectively dead for our needs and we needed to

switch to something else—our needs had outlived the capabilities of the external library. At the time of this writing, OpenStack projects are migrating towards using [alembic](#) instead, a new SQL database migrations tool with Python 3 support. This is happening not without some effort, but fortunately without much pain. OpenStack also had to adopt `sqlalchemy-migrate` in the meantime so it would be able to fix some of its bugs. That means even more work for the OpenStack community.

2.3.1 The External Libraries Safety Checklist

All of this builds up to one important question: how can you be sure you won't fall into this external libraries trap? Unfortunately, you can't: programmers are people, too, and there's no way you can know for sure whether a library that's zealously maintained today will still be in good shape in a few months. However, using such libraries may be worth the risk,; it's just important to carefully assess your situation. The following checklist helps when choosing whether to use an external library.

- **Python 3 compatibility** – There are still libraries that only supports only Python 2, and that's probably not a good sign of health nor a good option if you write your code using Python 3.
- **Active development** – [GitHub](#) and [Ohloh](#) usually provide enough information to determine whether a given library is still being worked on by its maintainers.
- **Active maintenance** – Even if a library is "finished" (i.e. feature-complete), the maintainers should be ensuring it remains bug-free. Check the project's tracking system to see how quickly the maintainers respond to bugs.
- **Packaged with OS distributions** – If a library is packaged with major Linux distributions, that means other projects are depending on it – so if something goes wrong, you won't be the only one complaining. It's also a good idea to check this if you plan to release your software to the public: your code will be easier to distribute if its dependencies are already installed on the end user's machine.

- **API compatibility commitment** – Nothing’s worse than having your software suddenly break because a library it depends on changed its entire API. You might want to check whether your chosen library has had anything like this happen in the past.
- **License** – You need to make sure that the license is compatible with the software you’re planning to write and that it allows you to do whatever you intend to do with your code in terms of distribution, modification, and execution.

Applying this checklist to dependencies is also a good idea, though that could turn out to be a huge undertaking. As a compromise, if you know your application is going to depend heavily on a particular library, you should apply this checklist to each of that library’s dependencies.

2.3.2 Protecting Your Code with an API Wrapper

No matter what libraries you end up using, you need to treat them as you would any other tools: like useful devices that could potentially do some serious damage. Therefore, for safety, libraries should be treated like any physical tool: kept in your tool shed, away from your fragile valuables but available when you actually need them.

No matter how useful an external library might be, you need to be wary of letting it get its hooks into your actual source code. Otherwise, if something goes wrong and you need to switch libraries, you might have to rewrite huge swaths of your program. A better idea is to write your own API — a wrapper that encapsulates your external libraries and keeps them out of your source code. Your program never has to know what external libraries it’s using; only what functionality your API provides. Then, if you need to use a different library, all you have to change is your wrapper: as long as the new library provides the same functionality, you won’t have to touch the rest of your codebase at all. There might be exceptions, but probably not many: most

libraries are designed to solve a tightly focused range of problems and can therefore be easily isolated.

Later, in Section 5.5.3, we'll also look at how you can use entry points to build driver systems that will allow you to treat parts of your projects as modules you can switch out at will.

2.4 Package Installation: Getting More From pip

The *pip* project offers a really simple way to handle package and external library installations. It is actively developed, well maintained, and included with Python starting at version 3.4. It can install or uninstall packages from the Python Packaging Index (PyPI), a tarball, or a *Wheel* archive (we'll discuss these in Section 5.3).

Its usage is simple:

```
1 $ pip install --user voluptuous
2 Downloading/unpacking voluptuous
3   Downloading voluptuous-0.8.3.tar.gz
4   Storing download in cache at ~/.cache/pip/https%3A%2F%2Fpypi.python. ←
   org%2Fpackages%2Fsource%2Fv%2Fvoluptuous%2Fvoluptuous-0.8.3.tar. ←
   gz
5   Running setup.py egg_info for package voluptuous
6
7 Requirement already satisfied (use --upgrade to upgrade): distribute ←
   in /usr/lib/python2.7/dist-packages (from voluptuous)
8 Installing collected packages: voluptuous
9   Running setup.py install for voluptuous
10
11 Successfully installed voluptuous
12 Cleaning up...
```

By looking it up on the [PyPI distribution index](#), PyPI is the *Python Package Index* where anyone can upload a package for distribution and installation by others.

2.4. PACKAGE INSTALLATION: GETTING MORE FROM PIP

You can also provide a `--user` option that makes *pip* install the package in your home directory. This avoids polluting your operating system directories with packages installed system-wide.

You can list the packages that are currently installed by using the `pip freeze` command, like so:

```
1 $ pip freeze
2 Babel==1.3
3 Jinja2==2.7.1
4 commando=0.3.4
5 ...
```

Uninstalling packages is also supported by *pip*, using the `uninstall` command:

```
1 $ pip uninstall pika-pool
2 Uninstalling pika-pool-0.1.3:
3 /usr/local/lib/python2.7/site-packages/pika_pool-0.1.3.dist-info/ ↩
4 DESCRIPTION.rst
5 /usr/local/lib/python2.7/site-packages/pika_pool-0.1.3.dist-info/ ↩
6 INSTALLER
7 /usr/local/lib/python2.7/site-packages/pika_pool-0.1.3.dist-info/ ↩
8 METADATA
9 /usr/local/lib/python2.7/site-packages/pika_pool-0.1.3.dist-info/ ↩
10 RECORD
11 /usr/local/lib/python2.7/site-packages/pika_pool-0.1.3.dist-info/ ↩
12 WHEEL
13 /usr/local/lib/python2.7/site-packages/pika_pool-0.1.3.dist-info/ ↩
14 metadata.json
15 /usr/local/lib/python2.7/site-packages/pika_pool-0.1.3.dist-info/ ↩
16 top_level.txt
17 /usr/local/lib/python2.7/site-packages/pika_pool.py
18 /usr/local/lib/python2.7/site-packages/pika_pool.pyc
19 Proceed (y/n)? y
20 Successfully uninstalled pika-pool-0.1.3
```

One very valuable feature of `pip` is its ability to install a package without copying the package's file. The typical use case for this feature is when you're actively working

2.4. PACKAGE INSTALLATION: GETTING MORE FROM PIP

on a package and want to avoid the long and boring process of reinstalling it each time you need to test a change. This can be achieved by using the `-e <directory>` flag:

```
1 $ pip install -e .
2 Obtaining file:///Users/jd/Source/daiquiri
3 Installing collected packages: daiquiri
4   Running setup.py develop for daiquiri
5 Successfully installed daiquiri
```

Here, `pip` does not copy the files from the local source directory but places a special file called an egg-link in your distribution path. For example:

```
1 $ cat /usr/local/lib/python2.7/site-packages/daiquiri.egg-link
2 /Users/jd/Source/daiquiri
```

The egg-link file contains the path to add to `sys.path` to look for packages. The result can be easily checked by running the following command:

```
1 $ python -c "import sys; print('/Users/jd/Source/daiquiri' in sys.path ↵
2   )"
3 True
```

Another useful `pip` tool is the `-e` option of `pip install`, helpful for deploying code from repositories of various version control systems: `git`, `Mercurial`, `Subversion` and even `Bazaar` are supported. For example, you can install any library directly from a `git` repository by passing its address as a URL after the `-e` option:

```
1 $ pip install -e git+https://github.com/jd/daiquiri.git#egg=daiquiri
2 Obtaining daiquiri from git+https://github.com/jd/daiquiri.git#egg= ↵
3   daiquiri
4 Cloning https://github.com/jd/daiquiri.git to ./src/daiquiri
5 Installing collected packages: daiquiri
6   Running setup.py develop for daiquiri
7 Successfully installed daiquiri
```

For the installation to work correctly, you need to provide the package egg name by adding `#egg=` at the end of the URL. Then, *pip* just uses `git clone` to clone the

repository inside a `src/<eggname>` and creates an `egg-link` file pointing to that same cloned directory.

This mechanism is extremely handy when depending on unreleased version of libraries or when working in a continuous testing system. However, since there is no versioning behind, the `-e` option can also be very nasty. You cannot know in advance that the next commit in this remote repository is not going to break everything.

Finally, all other installation tools are being deprecated in favor of *pip*, so can confidently treat it as your one-stop shop for all your package management needs.

2.5 Using and Choosing Frameworks

Python has a variety of frameworks available for various kinds of applications: if you're writing a web application, you could use *Django*, *Pylons*, *TurboGears*, *Tornado*, *Zope*, or *Plone*; if you're looking for an event-driven framework, you could use *Twisted* or *Circuits*; and so on.

The main difference between frameworks and external libraries is that applications use frameworks by building on top of them: your code will extend the framework rather than vice versa. Unlike a library, which is basically an add-on you can bring in to give your code some extra *oomph*, a framework forms the *chassis* of your code: everything you do builds on that chassis in some way, which can be a double-edged sword. There are plenty of upsides to using frameworks, such as rapid prototyping and development, but there are also some noteworthy downsides, such as lock-in. You need to take these considerations into account when you decide whether to use a framework.

The recommended method for choosing a framework for a Python application is largely the same as the one described in Section 2.3.1 - which makes sense, as frameworks are distributed as bundles of Python libraries. Sometimes frameworks also include tools for creating, running, and deploying applications, but that doesn't

change the criteria you should apply. We've already established that replacing an external library after you've already written code that makes use of it is a pain, but replacing a framework is a thousand times worse, usually requiring a complete rewrite of your program from the ground up.

To give an example, the Twisted framework mentioned earlier still doesn't have full Python 3 support: if you wrote a program using Twisted a few years back and wanted to update it to run on Python 3, you'd be out of luck. Either you'd have to rewrite your entire program to use a different framework, or you'd have to wait until someone finally gets around to upgrading it with full Python 3 support.

Some frameworks are lighter than others. For example, Django has its own built-in ORM functionality; Flask, on the other hand, has nothing of the sort. The *less* a framework tries to do for you, the fewer problems you'll have with it in the future. However, each feature a framework lacks is another problem for you to solve, either by writing your own code or going through the hassle of hand-picking another library to handle it. It's your choice which scenario you'd rather deal with, but choose wisely: migrating away from a framework when things go sour can be a Herculean task, and even with all its other features, there's nothing in Python that can help you with that.

2.6 Doug Hellmann on Python Libraries

Doug Hellmann is a senior developer at DreamHost and a fellow contributor to the OpenStack project. He launched the website [Python Module of the Week](#) and has written an excellent book called *The Python Standard Library by Example*. He is also a Python core developer. I've asked Doug a few questions about the Standard Library and designing libraries and applications around it.

“



When you start writing a Python application from scratch, what’s your first move? Is it different from hacking an existing application?

The steps are similar in the abstract, but the details change. There tend to be more differences between my approach to working on applications and libraries than there are for new versus existing projects.

When I want to change existing code, especially when it has been created by someone else, I start by digging in to figure out how it works and where my change would need to go. I may add logging or print statements, or use *pdb*, and run the application with test data to make sure I understand what it is doing. I usually make the change and test it by hand, then add any automated tests before contributing a patch.

I take the same exploratory approach when I create a new application. I create some code and run it by hand, then write tests to make sure I’ve covered all of the edge cases after I have the basic aspect of a feature working. Creating the tests may also lead to some refactoring to make the code easier to work with.

That was definitely the case with *smiley*. I started by experimenting with Python’s trace API using some throw-away scripts, before building the real application. My original vision for *smiley* included one

piece to instrument and collect data from another running application, and a second piece to collect the data sent over the network and save it. In the course of adding a couple of different reporting features, I realized that the processing for replaying the data that had been collected was almost identical to the processing for collecting it in the first place. I refactored a few classes, and was able to create a base class for the data collection, database access, and report generator. Making those classes conform to the same API allowed me to easily create a version of the data collection app that wrote directly to the database instead of sending information over the network.

While designing an app, I think about how the user interface works, but for libraries, I focus on how a developer will use the API. Thinking about how to write programs with the new library can be made easier by writing the tests first, instead of after the library code. I usually create a series of example programs in the form of tests, and then build the library to work that way.

I have also found that writing the documentation for a library before writing any code at all gives me a way to think through the features and workflows for using it without committing to the implementation details. It also lets me record the choices I made in the design so the reader understands not just how to use the library but the expectations I had while creating it. That was the approach I took with *stevedore*.

I knew I wanted *stevedore* to provide a set of classes for managing plugins for applications. During the design phase, I spent some time thinking about common patterns I had seen for consuming plugins and wrote a few pages of rough documentation describing how the classes would be used. I realized that if I put most of the complex arguments into the class constructors, the `map()` methods could be al-

most interchangeable. Those design notes fed directly into the introduction for stevedore’s official documentation, explaining the various patterns and guidelines for using plugins in an application.

What’s the process for getting a module into the Python Standard Library?

The full process and guidelines can be found in the [Python Developer’s Guide](#).

Before a module can be added to the Python Standard Library, it needs to be proven to be stable and widely useful. The module should provide something that is either hard to implement correctly or so useful that many developers have created their own variations. The API should be clear and the implementation should not have dependencies on modules outside the Standard Library.

The first step to proposing a new module is bringing it up within the community via the *python-ideas* list to informally gauge the level of interest. Assuming the response is positive, the next step is to create a Python Enhancement Proposal (PEP), which includes the motivation for adding the module and some implementation details of how the transition will happen.

Because package management and discovery tools have become so reliable, especially *pip* and the Python Package Index (PyPI), it may be more practical to maintain a new library outside of the Python Standard Library. A separate release allows for more frequent updates with new features and bugfixes, which can be especially important for libraries addressing new technologies or APIs.

What are the top three modules from the Standard Library that you wish people knew more about and would start using?

I've been doing a lot of work with dynamically loaded extensions for applications recently. I use the `abc` module to define the APIs for those extensions as abstract base classes to help extension authors understand which methods of the API are required and which are optional. Abstract base classes are built into some other object-oriented programming languages, but I've found a lot of Python programmers don't know we have them as well.

The binary search algorithm in the `bisect` module is a good example of a feature that is widely useful and often implemented incorrectly, which makes it a great fit for the Standard Library. I especially like the fact that it can search sparse lists where the search value may not be included in the data.

There are some useful data structures in the `collections` module that aren't used as often as they could be. I like to use `namedtuple` for creating small class-like data structures that just need to hold data but don't have any associated logic. It's very easy to convert from a `namedtuple` to a regular class if logic does need to be added later, since `namedtuple` supports accessing attributes by name. Another interesting data structure is `ChainMap`, which makes a good stackable namespace. `ChainMap` can be used to create contexts for rendering templates or managing configuration settings from different sources with clearly defined precedence.

A lot of projects, including OpenStack, or external libraries, roll their own abstractions on top of the Standard Library. I'm particularly thinking about things like date/time handling, for example. What would be your advice on that? Should programmers stick to the Standard Library, roll their own functions, switch to some external library, or start sending patches to Python?

All of the above! I prefer to avoid reinventing the wheel, so I advocate strongly for contributing fixes and enhancements upstream to projects that can be used as dependencies. On the other hand, sometimes it makes sense to create another abstraction and maintain that code separately, either within an application or as a new library.

The example you raise, the `timeutils` module in OpenStack, is a fairly thin wrapper around Python's `datetime` module. Most of the functions are short and simple, but by creating a module with the most common operations, we can ensure they are handled consistently throughout all OpenStack projects. Because a lot of the functions are application-specific, in the sense that they enforce decisions about things like timestamp format strings or what "now" means, they are not good candidates for patches to Python's library or to be released as a general purpose library and adopted by other projects.

In contrast, I have been working to move the API services in OpenStack away from the Web Server Gateway Interface (WSGI) framework created in the early days of the project and onto a third-party web development framework. There are a lot of options for creating WSGI applications in Python, and while we may need to enhance one to make it completely suitable for OpenStack's API servers, contributing those reusable changes upstream is preferable to maintaining a "private" framework.

Do you have any particular recommendations on what to do when importing and using a lot of modules, from the Standard Library or elsewhere?

I don't have a hard limit, but if I have more than a handful of imports, I reconsider the design of the module and think about splitting it up into a package. The split may happen sooner for a lower level module than for a high-level or application module, since at a higher level I

expect to be joining more pieces together.

Regarding Python 3, what are the modules that are worth mentioning and might make developers more interested in looking into it?

The number of third-party libraries supporting Python 3 has reached critical mass. It's easier than ever to build new libraries and applications for Python 3, and maintaining support for Python 2.7 is also easier thanks to the compatibility features added to 3.3. The major Linux distributions are working on shipping releases with Python 3 installed by default. Anyone starting a new project in Python should look seriously at Python 3 unless they have a dependency that hasn't been ported. At this point, though, libraries that don't run on Python 3 could almost be classified as "unmaintained."

Many developers write all their code into an application, but there are cases where it would be worth the effort to branch their code out into a Python library. In term of design, planning ahead, migration, etc., what are the best ways to do this?

Applications are collections of "glue code" holding libraries together for a specific purpose. Designing based on implementing those features as a library first and then building the application ensures that code is properly organized into logical units, which in turn makes testing simpler. It also means the features of an application are accessible through the library and can be remixed to create other applications. Failing to take this approach means the features of the application are tightly bound to the user interface, which makes them harder to modify and reuse.

What advice would you give to people planning to start their own Python libraries?

I always recommend designing libraries and APIs from the top down, applying design criteria such as the **Single Responsibility Principle (SRP)** at each layer. Think about what the caller will want to do with the library, and create an API that supports those features. Think about what values can be stored in an instance and used by the methods versus what needs to be passed to each method every time. Finally, think about the implementation and whether the underlying code should be organized differently from the public API.

SQLAlchemy is an excellent example of applying those guidelines. The declarative object relational mapping (ORM), data mapping, and expression generation layers are all separate. A developer can decide the right level of abstraction for entering the API and using the library based on their needs rather than constraints imposed by the library's design.

What are the most common programming errors that you encounter while reading random Python developers' code?

A big area where Python's idioms are different from other languages is looping and iteration. For example, one of the most common anti-patterns I see is using a for loop to filter one list by appending items to a new list and then processing the result in a second loop (possibly after passing the list as an argument to a function). I almost always suggest converting filtering loops like that to generator expressions because they are more efficient and easier to understand. It's also common to see lists being combined so their contents can be processed together in some way, rather than using `itertools.chain()`.

There are also some more subtle things I suggest in code reviews, like using a `dict()` as a lookup table instead of a long `if: then: else` block; making sure functions always return the same type of object (e.g., an empty list instead of `None`); reducing the number of arguments to a

function by combining related values into an object with either a tuple or a new class; and defining classes to use in public APIs instead of relying on dictionaries.

Do you have a concrete example, something you’ve either done or witnessed, of picking up a “wrong” dependency?

Recently, I had a case in which a new release of *pyarsing* dropped Python 2 support and caused me a little trouble with a library I maintain. The update to *pyarsing* was a major revision, and was clearly labeled as such, but because I had not constrained the version of the dependency in the settings for *cliff*, the new release of *pyarsing* caused issues for some of *cliff*’s consumers. The solution was to provide different version bounds for Python 2 and Python 3 in the dependency list for *cliff*. This situation highlighted the importance of both understanding dependency management and ensuring proper test configurations for continuous integration testing.

What’s your take on frameworks?

Frameworks are like any other kind of tool. They can help, but you need to take care when choosing one to make sure that it’s right for the job at hand.

By pulling out the common parts into a framework, you can focus your development efforts on the unique aspects of an application. They also help you bring an application to a useful state more quickly than if you started from scratch by providing a lot of bootstrapping code for doing things like running in development mode and writing a test suite. They also encourage you to be consistent in the way you implement the application, which means you end up with code that is easier to understand and more reusable.

There are some potential pitfalls to watch out for when working with frameworks, though. The decision to use a particular framework usu-

ally implies something about the design of the application itself. Selecting the wrong framework can make an application harder to implement if those design constraints do not align naturally with the application's requirements. You may end up fighting with the framework if you try to use different patterns or idioms than it recommends.

”

CHAPTER 3



Documentation and Good API Practice

In this chapter, we'll discuss documentation; specifically, how to automate the trickier and more tedious aspects of documenting your project with Sphinx.

While you will still have to write the documentation yourself, Sphinx will simplify your task. As it is common to provide features using a Python library, we'll also look at how to manage and document your public API changes.

Because your API will have to evolve as you add, remove, and change the features it provides, it's rare to get everything built perfectly from the outset, but I'll show you a few things you can do to ensure your API is as user-friendly as possible. We'll end this chapter with an interview with Christophe de Vienne, author of the Web Services Made Easy framework, in which he discusses best practices for developing and maintaining APIs.

3.1 Documenting with Sphinx

Documentation is one of the most important parts of writing software. Unfortunately, a lot of projects don't provide proper documentation. Writing documentation is seen as a complicated and daunting task, but it doesn't have to be: with the tools available to Python programmers, documenting your code can be just as easy as writing it in the first place.

One of the biggest reasons for sparse or nonexistent documentation is that many people assume the only way to document code is by hand. Even with multiple people on a project, this means one or more of your team will end up having to juggle contributing code with maintaining documentation — and if you ask any developer which job they'd prefer, you can be sure they'll say they'd rather write software than write about software.

Sometimes the documentation process is even completely separate from the development process, meaning that the documentation is written by people who have never written so much as a line of the actual code. Furthermore, any documentation produced this way is likely to be out-of-date: whether the documentation is handled by the programmers themselves or by dedicated writers, it's almost impossible for manual documentation to keep up with the pace of development.

Here's the bottom line: the more degrees of separation between your code and your documentation, the harder it will be to keep the latter properly maintained. So why keep your code and documentation separate at all? It's not only possible to put your documentation directly in the code itself, but it's also simple to convert that documentation into easy-to-read HTML and PDF files.

The de facto standard documentation format for Python is *reStructuredText*, or *reST* for short. It's a lightweight markup language (like the famous *Markdown*) that's as easy to read and write for humans as it is for computers. Sphinx is the most commonly used tool for working with this format; Sphinx can read reST-formatted content and output documentation in a variety of other formats.

I recommend that your project documentation always include the following:

- The problem your project is intended to solve, in one or two sentences.
- The license your project is distributed under. If your software is open source, you should also include this information in a header in each code file; just because you've uploaded your code to the Internet doesn't mean that people will know what they're allowed to do with it.
- A small example of how your code works.
- Installation instructions.
- Links to community support, mailing list, IRC, forums, and so on.
- A link to your bug tracker system.
- A link to your source code so that developers can download and start delving into it right away.

You should also include a `README.rst` file that explains what your project does. This README will be displayed on your [GitHub](#) or [PyPI](#) project page; both sites know how to handle *reST* formatting.



Tip

If you're using [GitHub](#), you can also add a `CONTRIBUTING.rst` file that will be displayed when someone creates a pull request. It should provide a checklist for them to follow before they submit the request, e.g. follow PEP 8 or don't forget to run the unit tests.

Tip

Read the Docs allows you to build and publish your documentation online automatically. Signing up and configuring a project is straightforward. Then Read the Docs searches for your *Sphinx* configuration file, builds your documentation, and makes it available for your users to access. It's a great companion to code hosting sites.

3.2 Getting started with Sphinx and reST

You can get Sphinx from [its web site](#). There are instructions on installing Sphinx on the site, but the easiest method is to install with `pip install sphinx`.

Once Sphinx is installed, run `sphinx-quickstart` in your project's top-level directory. This will create the directory structure that Sphinx expects to find, along with two files in the `doc/source` folder: `conf.py`, which contains Sphinx's configuration settings (and is absolutely required for Sphinx to work), and `index.rst`, which serve as the front page of your documentation. Once you run the quick-start command, you'll be taken through a series of steps to designate naming conventions, version conventions, and options for many other useful tools and standards.

The `conf.py` file contains a few documented variables, such as the project name, the author, and the theme to use for HTML output. Feel free to edit this file at your convenience.

Once you've built your structure and set your defaults, you can build your documentation in HTML by calling `sphinx-build` with your source directory and output directory as arguments, as shown in Example 3.1. The command `sphinx-build` reads the `conf.py` file from the source directory and parses all the `.rst` files from this directory. It then renders them in HTML in the output directory.

Example 3.1 Building a basic Sphinx HTML document

```
1 $ sphinx-build doc/source doc/build
2   import pkg_resources
3 Running Sphinx v1.2b1
4 loading pickled environment... done
5 No builder selected, using default: html
6 building [html]: targets for 1 source files that are out of date
7 updating environment: 0 added, 0 changed, 0 removed
8 looking for now-outdated files... none found
9 preparing documents... done
10 writing output... [100%] index
11 writing additional files... genindex search
12 copying static files... done
13 dumping search index... done
14 dumping object inventory... done
15 build succeeded.
```

Now you can open `doc/build/index.html` in your favorite browser and read your documentation.

**Tip**

If you're using *setuptools* for packaging, *Sphinx* extends it to support the command `setup.py build_sphinx`, which will run `sphinx-build` automatically.

Your documentation begins with the `index.rst` file, but it doesn't have to end there: reST supports include directives to include reST files from other reST files, so there's nothing stopping you from dividing your documentation up into multiple files. Don't worry too much about syntax and semantics to start with: it's true that reST offers a lot of formatting possibilities, but you'll have plenty of time to dive into the reference later. The [complete](#) reference explains how to create titles, bulleted lists, tables, and more.

3.3 Sphinx Modules

Sphinx is highly extensible: its basic functionality only supports manual documentation, but it comes with a number of useful modules which enable automatic documentation and other features. For example, `sphinx.ext.autodoc` extracts reST-formatted docstrings from your modules and generates `.rst` files for inclusion. This is one of the options `sphinx-quickstart` will ask you if you want to activate. If you didn't select that option, you can still edit your `conf.py` file and add it as an extension like so:

```
1 extensions = ['sphinx.ext.autodoc']
```

Note that `autodoc` will **not** automatically recognize and include your modules. You need to explicitly indicate which modules you want to be documented by adding something like Example 3.2 to one of your `.rst` files.

Example 3.2 Indicating the modules for autodoc to document

```
1 .. automodule:: foobar
2    :members: ❶
3    :undoc-members: ❷
4    :show-inheritance: ❸
```

- ❶ Request all documented members to be printed (optional)
- ❷ Request all undocumented members to be printed (optional)
- ❸ Show inheritance (optional)

Also note the following:

- If you don't include any directives, Sphinx won't generate any output.
- If you only specify `:members:`, undocumented nodes on your module/class/method tree will be skipped, even if all their members are documented. For example, if you

3.4. AUTOMATING THE TABLE OF CONTENTS WITH AUTOSUMMARY

document the methods of a class but not the class itself, `:members:` will exclude both the class and its methods entirely. To keep this from happening, you'd either have to write a docstring for the class or specify `:undoc-members:` as well.

- Your module needs to be where Python can import it. Adding `.`, `..`, and/or `../..` to `sys.path` can help with this.

The *autodoc* gives you the power to include most of your documentation in your actual source code. You can even pick and choose which modules and methods to document – it's not an "all-or-nothing" solution. By maintaining your documentation directly alongside your source code, you can easily ensure it stays up-to-date.

3.4 Automating the Table of Contents with autosummary

If you're writing a Python library, you'll usually want to format your API documentation with a table of contents containing links to individual pages for each module.

The `sphinx.ext.autosummary` module was created specifically to handle this common use case. First, you need to enable it in `conf.py` by adding the following line:

```
1 extensions = ['sphinx.ext.autosummary']
```

Then, you can add something like the following to an `.rst` file to automatically generate a table of contents for the specified modules:

```
1 .. autosummary::
2
3   mymodule
4   mymodule.submodule
```

This will create files called `generated/mymodule.rst` and `generated/mymodule.submodule.rst` containing the autodoc directives described earlier. Using this same format, you can specify which parts of your module API you want included in your documentation.



Tip

The `sphinx-apidoc` command can automatically create these files for you; check out the Sphinx documentation to find out more.

3.5 Automating Testing with doctest

Another useful feature of *Sphinx* is the ability to run doctest on your examples automatically when you build your documentation. The standard Python doctest module searches your documentation for code snippets and runs them to test whether they accurately reflect what your code does. Every paragraph starting with the primary prompt `>>>` is treated as a code snippet to test. For example, if you wanted to document the standard `print` function from Python, you could write this documentation snippet and doctest would make sure that the result was what it should be:

```
1 To print something to the standard output, use the :py:func:`print` ↵  
  function.  
2  
3 >>> print("foobar")  
4 foobar
```

Having such examples in your documentation is a good way for users to understand your API. However, it's easy to put off and eventually forget to update your examples as your API evolves. Fortunately, doctest helps make sure this doesn't hap-

3.5. AUTOMATING TESTING WITH DOCTEST

pen. If your documentation includes a step-by-step tutorial, `doctest` will help you keep it up-to-date throughout development by testing every line it can.

You can also use `doctest` for documentation-driven development (DDD): write your documentation and examples first and then write code to match your documentation. Taking advantage of this feature is as simple as running `sphinx-build` with the special `doctest` builder, like this:

```
1  $ sphinx-build -b doctest doc/source doc/build
2  Running Sphinx v1.2b1
3  loading pickled environment... done
4  building [doctest]: targets for 1 source files that are out of ←
   date
5  updating environment: 0 added, 0 changed, 0 removed
6  looking for now-outdated files... none found
7  running tests...
8
9  Document: index
10 -----
11  1 items passed all tests:
12    1 tests in default
13  1 tests in 1 items.
14  1 passed and 0 failed.
15  Test passed.
16
17  Doctest summary
18  =====
19    1 test
20    0 failures in tests
21    0 failures in setup code
22    0 failures in cleanup code
23  build succeeded.
```

When using the `doctest` builder, *Sphinx* reads the usual `.rst` files and executes code examples that are contained in those files.

Sphinx also provides a bevy of other features, either out of the box or through

extension modules, including these:

- Link between projects
- HTML themes
- Diagrams and formulas
- Output to Texinfo and EPUB format
- Linking to external documentation

You might not need all this functionality right away, but if you ever need it in the future, it's good to know in advance that there are modules that can provide it. Again, check out the full Sphinx documentation to find out more.

3.6 Writing a Sphinx Extension

Sometimes the off-the-shelf solutions just aren't enough and you need to create custom tools to deal with a situation.

Say you're writing an HTTP REST API. Sphinx will only document the Python side of your API, forcing you to write your reST API documentation by hand, with all the problems that entails. The creators of Web Services Made Easy (WSME) (interviewed at the end of this chapter) have come up with a solution: a Sphinx extension called `sphinxcontrib-pecanwsme` that analyzes docstrings and actual Python code to generate reST API documentation automatically.



Tip

For other HTTP frameworks, such as Flask, Bottle, and Tornado, you can use [*sphinxcontrib.httpdomain*](#).

My point here is that whenever you know you could extract information from your code that could help to build documentation, you should, and you should also automate the process. This is better than trying to maintain manually written documentation, especially when you can leverage auto-publication tools such as [Read the Docs](#).

We'll examine the `sphinxcontrib-pecanwsme` extension as an example of writing your own Sphinx extension. The first step is to write a module — preferably as a submodule of `sphinxcontrib`, as long as your module is generic enough — and pick a name for it. Sphinx requires this module to have one predefined function called `setup(app)`, which contains the methods you'll use to connect your code to Sphinx events and directives. The full list of methods is available in the [Sphinx extension API documentation](#).

For example, the `sphinxcontrib-pecanwsme` extension includes a single directive called `rest-controller`, added using the `setup(app)` function. This added directive needs a fully qualified controller class name to generate documentation for, as shown in Example 3.3.

Example 3.3 Code from `sphinxcontrib.pecanwsme.rest.setup`

```
1 def setup(app):  
2     app.add_directive('rest-controller', RESTControllerDirective)
```

The `add_directive` method in Example 3.3 registers the directive `rest-controller` and delegates its handling to the `RESTControllerDirective` class. This `RESTControllerDirective` class exposes certain attributes that indicate how the directive treats content, whether it has arguments, and so on. The class also implements a `run()` method that does the actual work of extracting the documentation from your code and returns parsed data to *Sphinx*. The [sphinx-contrib](#) repository at [has a bunch of small modules](#) that can help you develop your own extensions.

**Note**

Even though *Sphinx* is written in Python and targets it by default, extensions are available that allow it to support other languages as well. You can use *Sphinx* to document your project in full, even if it uses multiple languages at once.

As another example, in one of my projects named **Gnocchi** — a database for storing and indexing time series data at a large scale — I’ve used a custom Sphinx extension to autogenerate documentation. Gnocchi provides a REST API, and usually to document such an API, projects will manually write examples of what an API request should look like and what the response would look like. Unfortunately, this approach is error prone and always out of sync with reality.

Using the unit-testing code available to test the Gnocchi API, we built a Sphinx extension to run Gnocchi and generate an `.rst` file containing HTTP requests and responses run against a real Gnocchi server. In this way, we ensure the documentation is up-to-date: the server responses are not manually crafted, and if a manually written request fails, then the documentation process fails, and we know that we must fix the documentation.

Including that code in the book would be too verbose, but you can check the sources of Gnocchi online and take a look at the `gnocchi.gendoc` module to get an idea of how it works.

3.7 Managing Changes to Your API

Well-documented code is a sign to other developers that the code is good for being imported and used to build something else. When building a library and exporting an API for other developers to use, for example, you want to provide the reassurance of solid documentation.

This section will deal with good practice for public APIs. These will be exposed to users of your library or application, and while you can do whatever you like with internal APIs, public APIs should be handled with care.

To distinguish between public and private APIs, the Python convention is to prefix the symbol for a private API with an underscore: `foo` is public, but `_bar` is private. You should use this convention both to recognize whether another API is public or private and to name your own APIs.

In contrast to other languages, such as Java, Python does not enforce any restriction on accessing code marked as private or public. The naming conventions are just to facilitate understanding among programmers.

3.7.1 Numbering API Versions

When properly constructed, the version number of an API can give user a great deal of information. Python has no particular system or convention in place for numbering API versions, but we can take inspiration from Unix platforms, which use a complex management system for libraries with fine-grained version identifiers.

Generally, your version numbering should reflect changes in the API that will impact users; for example, use major number increments for major changes and more granular increments for minor updates. A good example to look at is the [Python requests library](#). This library increments its API numbers based on the number of changes made with each new version and the impact the changes might have on consuming programs. When the API had a major change, the major version number changed from 1 to 2. When only a few new API calls are added, the lesser number is bumped, so the version might go from 2.2 to 2.3. If a change only involves bug fixes, the version is bumped from 2.2.0 to 2.2.1, for example.

Version numbers hint to developers that they should look at changes between two releases of a library, but alone they are not enough to fully guide a developer: you must provide detailed documentation to describe those changes.

3.7.2 Documenting Your API Changes

Whenever you make changes to an API, the first and most important thing to do is to heavily document the change. Without documentation, it is impossible for a consumer of your code to get a quick overview of what's new and what's going away. Your document should cover the following:

- New elements of the new interface
- Elements of the old interface that are deprecated
- Instructions on how to migrate to the new interface

You should also make sure that you don't remove the old interface right away. I recommend keeping the old interface until it becomes too much trouble to do so. If you have marked it as deprecated, users will know not to use it.

As an example of an API with well-documented changes, Example 3.4 is an example of API change documentation for code that provides a representation of a car object that can turn in any direction. For whatever reason, the developers decided to retract the `turn_left` method and instead provide a generic `turn` method that can take the direction as an argument.

Example 3.4 An example of API change documentation for a car object

```

1 class Car(object):
2     def turn_left(self):
3         """Turn the car left.
4
5         .. deprecated:: 1.1
6             Use :func:`turn` instead with the direction argument set to ←
7             left
8         """
9         self.turn(direction='left')
```

3.7. MANAGING CHANGES TO YOUR API

```
10 def turn(self, direction):
11     """Turn the car in some direction.
12
13     :param direction: The direction to turn to.
14     :type direction: str
15     """
16     # Write actual code for the turn function here instead
17     pass
```

The triple quotes here, `"""`, indicate the start and end of the docstrings, which will be pulled into the documentation when the user enters `help(Car.turn_left)` into the terminal or extracts the documentation with an external tool such as *Sphinx*. The deprecation of the `car.turn_left` method is indicated by `..deprecated 1.1`, where 1.1 refers to the first version released that ships this code as deprecated.

By using this deprecation method and making it visible via *Sphinx*, you'll make it clear to users that the function should not be used, and they will receive direct access to the new function along with an explanation of how to migrate old code.

Figure 3.1 shows *Sphinx* documentation that explains some deprecated functions.

complete

Returns a bool indicating whether the current frame is complete or not. If the frame is complete then `frame_size` will not increment any further, and will reset for the next frame.

Changed in version 1.5: Deprecated `header` and `keyframe` attributes and added the new `frame_type` attribute instead.

Changed in version 1.9: Added the `complete` attribute.

header

Contains a bool indicating whether the current frame is actually an SPS/PPS header. Typically it is best to split an H.264 stream so that it starts with an SPS/PPS header.

Deprecated since version 1.5: Please compare `frame_type` to `PiVideoFrameType.sps_header` instead.

Figure 3.1: Explanation of some deprecated functions

The downside of this approach is that it relies on developers reading your changelog or documentation when they upgrade to a newer version of your Python package. However, there is a solution for that: mark your deprecated functions with the `warnings` module.

3.7.3 Marking Deprecated Functions with the `warnings` Module

Though deprecated modules should be marked well enough in documentation that users will not attempt to call them, Python also provides the `warnings` module, which allows your code to issue various kinds of warnings when a deprecated function is called. These warnings, `DeprecationWarning` and `PendingDeprecationWarning`, can be used to tell the developer that a function they're calling is deprecated or going to

be deprecated, respectively. In this way, they'll be aware that they're using an old interface and should switch.



Note

For those who work with C, this is a handy counterpart to the `__attribute__((deprecated))` GCC extension.

To go back to the car object example in Example 3.4, we can make use of this to warn users when they are attempting to call deprecated functions, as shown in Example 3.5.

Example 3.5 A documented change to the car object API using the warning module

```

1 import warnings
2
3 class Car(object):
4     def turn_left(self):
5         """Turn the car left.
6
7         .. deprecated:: 1.1
8             Use :func:`turn` instead with the direction argument set to ←
9             "left".
10        """
11        warnings.warn("turn_left is deprecated, use turn instead",
12                      DeprecationWarning)
13        self.turn(direction='left')
14
15    def turn(self, direction):
16        """Turn the car in some direction.
17
18        :param direction: The direction to turn to.
19        :type direction: str
20        """
21        # Write actual code here instead

```

pass

Here, the `turn_left` function has been deprecated. By adding the `warnings.warn` line, we can write our own error message. Now, if any code should call the `turn_left` function, a warning will be issued that looks like this:

```
1 >>> Car().turn_left()
2 __main__:8: DeprecationWarning: turn_left is deprecated, use turn ←
   instead
```

Python 2.7 and later versions, by default, do not print any warnings emitted by the warning module because the warnings are filtered. To see those warnings printed, you need to pass the `-W` option to the Python executable. The option `-W all` will print all warnings to `stderr`. See the Python man page for more information on the possible values for `-W`.

When running test suites, developers can run Python with the `-W error` option, which will raise an error every time an obsolete function is called. Developers using your library can readily know exactly where their code needs to be fixed. Example 3.6 shows how Python transforms warnings into fatal exceptions when Python is called with the `-W error` option.

Example 3.6 Running python `-W error` and getting a deprecation error

```
1 >>> import warnings
2 >>> warnings.warn("This is deprecated", DeprecationWarning)
3 Traceback (most recent call last):
4   File "<stdin>", line 1, in <module>
5 DeprecationWarning: This is deprecated
```

Warnings are usually missed at runtime, and running a production system with the `-W error` option is rarely a good idea. However, running the test suite of a Python application with the `-W error` option can be a good way to catch warnings and fix them early on.

However, writing all those warnings, docstring updates, and so on manually can become tedious, so the `debtcollector` library has been created to help automate some of that. The `debtcollector` library provides a few decorators that you can use with your functions so that the correct warnings are emitted and the docstring is updated correctly. This way, you don't have to do these things manually. Example 3.7 shows how you can, with a simple decorator, indicate that a function has been moved to some other place.

Example 3.7 An API change automated with `debtcollector`

```

1 from debtcollector import moves
2
3 class Car(object):
4     @moves.moved_method('turn', version='1.1')
5     def turn_left(self):
6         """Turn the car left."""
7
8         return self.turn(direction='left')
9     def turn(self, direction):
10        """Turn the car in some direction.
11
12        :param direction: The direction to turn to.
13        :type direction: str
14        """
15        # Write actual code here instead
16        pass

```

Here we're using the `moves` method from `debtcollector`, whose `moved_method` decorator makes `turn_left` emit a `DeprecationWarning` whenever it's called.

3.8 Summary

Sphinx is the de facto standard for documenting Python projects. It supports a large variety of syntax, and it is easy to extend to add new syntax or new features if your

project has particular needs. Sphinx can also automate a variety of tasks, such as generating indexes or extracting documentation from your code, making it easy to maintain documentation in the long run.

Documenting changes to your API is critical, especially when you deprecate functionality, so that users are not caught unawares. Ways to document deprecations include the Sphinx deprecated keyword and the warning module, and the `debtcollector` library can automate maintaining this documentation.

3.9 Interview with Christophe de Vienne

Christophe de Vienne is a Python developer and the author of the *WSME* (Web Services Made Easy) framework, which allows developers to define web services in a Pythonic way and supports a wide variety of APIs, allowing it to be plugged into many other web frameworks.

“



What are the mistakes developers often make when designing a Python API?

There are a few common mistakes I avoid when designing a Python API by following these rules:

- **Don't make it too complicated.** Keep it simple. Complicated APIs are hard to understand and hard to document. While the actual library functionality doesn't have to be simple as well, it's smart to make it simple so users can't easily make mistakes. For example, the library is very simple and intuitive, but it does complex things behind the scenes. The `urllib` API, by contrast, is almost as complicated as the things it does, making it hard to use.
- **Make the magic visible.** When your API does things that your documentation doesn't explain, your end users will want to crack open your code and see what's going on under the hood. It's okay if you've got some magic happening behind the scenes, but your end users should never see anything unexpected happening up front, or they could become confused or rely on a behavior that may change.
- **Don't forget use cases.** When you're so focused on writing code, it's easy to forget to think about how your library will actually be used. Thinking up good use cases makes it easier to design an API.
- **Write unit tests.** TDD (test-driven development) is a very efficient way to write libraries, especially in Python, because it forces the developer to assume the role of the end user from the very beginning, which leads the developer to design for usability. It's the only approach I know of that allows a programmer to completely rewrite a library, as a last resort.

What aspects of Python may affect how easy it is to design a library API?

Python has no built-in way to define which sections of the API are public and which are private, which can be both a problem and an advantage. It's a problem because it can lead developer to not think as much as they should about which parts of their API are public and which parts should remain private. But with a little discipline,

documentation, and (if needed) tools like `zope.interface`, it doesn't stay a problem for long. It's an advantage when it makes it quicker and easier to refactor APIs while keeping compatibility with previous versions.

What do you consider when thinking about your API's evolution, deprecation, and removal?

There are several criteria I weigh when making any decision regarding API development:

- **How difficult will it be for users of the library to adapt their code?** Considering that there are people relying on your API, any change you make has to be worth the effort needed to adopt it. This rule is intended to prevent incompatible changes to the parts of the API that are in common use. That said, one of the advantages of Python is that it's relatively easy to refactor code to adopt an API change.
- **How easy will it be to maintain my API?** Simplifying the implementation, cleaning up the codebase, making the API easier to use, having more complete unit tests, making the API easier to understand at first glance... all of these things will make your life as a maintainer easier.
- **How can I keep my API consistent when applying a change?** If all the functions in your API follow a similar pattern (such as requiring the same parameter in the first position), make sure new functions follow that pattern as well. Also, doing too many things at once is a great way to end up doing none of them right: keep your API focused on what it's meant to do.
- **How will users benefit from the change?** Last but not least, always consider the users' point of view.

What advice do you have regarding API documentation in Python?

Good documentation makes it easy for newcomers to adopt your library. Neglecting it will drive away a lot of potential users — not just beginners, either. The problem is, documenting is difficult, so it gets neglected all the time!

- **Document early and include your documentation build in continuous integration.** With the Read the Docs tool for creating and hosting documentation, there's no excuse for not having documentation built and published (at least for open source software).
- Use docstrings to document classes and functions in your API. If you follow the [PEP 257](#) guidelines, developers won't have to read your source to understand what your API does. Generate HTML documentation from your docstrings — and don't limit it to the API reference.
- Give practical examples throughout. Have at least one “startup guide” that will show newcomers how to build a working example. The first page of the documentation should give a quick overview of your API's basic and representative use case.
- Document the evolution of your API in detail, version by version. Version control system (VCS) logs are not enough!
- Make your documentation accessible and, if possible, comfortable to read. Your users need to be able to find it easily and get the information they need without feeling like they're being tortured. Publishing your documentation through PyPI is one way to achieve this; publishing on Read the Docs is also a good idea, since users will expect to find your documentation there.

- Finally, choose a theme that is both efficient and attractive. I chose the “Cloud” Sphinx theme for WSME, but there are plenty of other themes out there to choose from. You don’t have to be a web expert to produce nice-looking documentation.

”

CHAPTER 4



Handling Timestamps and Time Zones

Time zones are complicated. Most people expect dealing with time zones to involve merely adding or subtracting a few hours from the universal time reference, UTC (Coordinated Universal Time), from -12 hours to +12 hours. However, reality shows otherwise: time zones are not logical or predictable. There are time zones with 15-minute granularity; countries that change time zones twice a year; countries that use a custom time zone during summer, known as daylight saving time, that starts on different dates; plus tons of special and corner cases. These make the history of time zones really interesting, but also complicate the handling of time zones in a logical or ordered way. All of those particularities should make you stop and think when dealing with time zones.

This chapter will outline why dealing with time zones is tricky and how you can best handle them in your programs. We'll look at how to build timestamp objects, how and why to make them time zone aware, and how to deal with some corner cases you might come across.

4.1 The Problem of Missing Time Zones

A timestamp without any time zone attached does not give any useful information, because without the necessary context given by the time zone, you cannot infer what point in time your application is really referring to. Without their respective time zones, therefore, you can't compare two timestamps; that would be like comparing days of the week without accompanying dates – whether Monday is before or after Tuesday depends on what weeks they're in. Timestamps without time zones attached should be considered irrelevant.

For that reason, your application should never have to handle timestamps with no time zone. Instead, it must raise an error if no time zone is provided, or it should make clear what default time zone is assumed – for example, it's common practice to choose for Coordinated Universal Time (UTC) as the default time zone.

You also must be careful of making any kind of time zone conversion before *storing* your timestamps. Imagine a user creates a recurring event every Wednesday at 10:00 AM in their local time zone, say Central European Time (CET). CET is an hour ahead of UTC, so if you convert that timestamp to UTC to store it, the event will be stored as every Wednesday at 09:00 AM. The CET time zone switches from UTC+01:00 to UTC+02:00 in the summer, so on top of that, in the summer months, your application will compute that the event starts at 11:00 AM CET every Wednesday. You can see how this program quickly becomes redundant!

Now that you understand the general problem of handling time zones, let's dig into our favorite language. Python comes with a timestamp object named `datetime.datetime` that can store date and time precise to the microsecond. The `datetime.datetime` object is qualified as being either time zone aware, in which case it embeds time zone information, or time zone unaware, in which case it does not. Unfortunately, the `datetime` API returns a time zone-unaware object by default, as you'll soon see in Example 4.1. Let's take a look at how to build a default timestamp object and then how to rectify it so that it uses time zones.

4.2 Building Default datetime Objects

To build a datetime object with the current date and time as values, you can use the `datetime.datetime.utcnow()` function. This function retrieves the date and time for the UTC time zone right now, as shown in Example 4.1. To build this same object using the date and time for the time zone of the region the machine is in, you can use the `datetime.datetime.now()` method. Example 4.1 retrieves the time and date for both UTC and my region's time zone.

Example 4.1 Getting the time of the day

```
1 >>> import datetime
2 >>> datetime.datetime.utcnow()
3 datetime.datetime(2018, 6, 15, 13, 24, 48, 27631)
4 >>> datetime.datetime.utcnow().tzinfo is None
5 True
```

We import the `datetime` library and define the `datetime` object as using the UTC time zone. This returns a UTC timestamp whose values are year, month, date, hours, minutes, seconds, and microseconds, respectively, in the listing. We can check if whether this object has time zone information by checking the `tzinfo` object, and here we're told that it doesn't.

We then create the `datetime` object using the `datetime.datetime.now()` method to retrieve the current date and time in the default time zone for the region of the machine:

```
1 >>> import datetime
2 >>> datetime.datetime.utcnow()
3 datetime.datetime(2018, 6, 13, 13, 27, 43, 24611)
```

This timestamp, too, is returned without any time zone, as we can tell from the absence of the `tzinfo` field – if the time zone information had been present, it would have appeared at the end of the output as something like `tzinfo=<UTC>`.

4.3. TIME ZONE-AWARE TIMESTAMPS WITH DATEUTIL

The datetime API always returns unaware datetime objects by default, and since there is no way for you to tell what the time zone is from the output, these objects are pretty useless.

Armin Ronacher, creator of the Flask framework, [proposes that an application always consider the unaware datetime objects from Python as UTC](#). However, as we just saw, that statement cannot be considered true for objects returned by `datetime.datetime.now()`, so I have different recommendations for building datetime objects, and the first is to always make sure they are time zone aware. That makes sure you can always compare your objects directly and check whether they are returned correctly, with the information you need. Let's see how to create time zone-aware timestamps using `tzinfo` objects.

You can also build your own datetime object with a particular date by passing the values you want for the different components of the day, as shown in Example 4.2.

Example 4.2 Constructing a datetime object from a date

```
1 >>> import datetime
2 >>> datetime.datetime(2018, 6, 15, 13, 24, 48, 27631)
3 datetime.datetime(2018, 6, 15, 13, 24, 48, 27631)
```

4.3 Time Zone-Aware Timestamps with dateutil

There are already many databases of existing time zones, maintained by central authorities such as IANA (Internet Assigned Numbers Authority), which are shipped with all major operating systems. For this reason, rather than creating our own time zone classes and manually duplicating those in each Python project, Python developers rely on the `dateutil` project to obtain `tzinfo` classes. The `dateutil` project provides the Python module `tz`, which makes time zone information available directly, without much effort: the `tz` module can access the operating system's time

4.3. TIME ZONE-AWARE TIMESTAMPS WITH DATEUTIL

zone information, as well as ship and embed the time zone database so it is directly accessible from Python.

You can install `dateutil` with the command `pip install python-dateutil`. The `dateutil` API allows you to obtain a `tzinfo` object based on a time zone name, like so:

```
1 >>> from dateutil import tz
2 >>> tz.gettz("Europe/Paris")
3 tzfile('/usr/share/zoneinfo/Europe/Paris')
4 >>> tz.gettz("GMT+1")
5 tzstr('GMT+1')
```

The `dateutil.tz.gettz` method returns an object implementing the `tzinfo` interface. This method accepts various string formats as argument, such as the time zone based on a location (for example, “Europe/Paris”) or a time zone relative to GMT. The `dateutil` time zone objects can be used as `tzinfo` classes directly, as demonstrated in Example 4.3.

Example 4.3 Using `dateutil` objects as `tzinfo` classes

```
1 >>> import datetime
2 >>> from dateutil import tz
3 >>> now = datetime.datetime.now()
4 >>> now
5 datetime.datetime(2018, 1, 16, 19, 40, 18, 234)
6 >>> tz = tz.gettz("Europe/Paris")
7 >>> now.replace(tzinfo=tz)
8 datetime.datetime(2017, 10, 16, 19, 40, 18, 279100, tzinfo=tzfile('/ ←
   usr/share/zoneinfo/Europe/Paris'))
```

As long as you know the name of the desired time zone, you can obtain a `tzinfo` object that matches the time zone you target. The `dateutil` module is able to access the time zone managed by the operating system, and if that information is for some reason unavailable, will fall back on its own list of embedded time zone. If you ever

4.3. TIME ZONE-AWARE TIMESTAMPS WITH DATEUTIL

need to access its embedded list of time zones, you can do so via the `dateutil.zoneinfo` module like so:

```
1 >>> from dateutil.zoneinfo import get_zonefile_instance
2 >>> zones = list(get_zonefile_instance().zones)
3 >>> sorted(zones)[:5]
4 ['Africa/Abidjan', 'Africa/Accra', 'Africa/Addis_Ababa', 'Africa/ ←
   Algiers', 'Africa/Asmara']
5 >>> len(zones)
6 592
```

In some cases, your program does not know which time zone it's running in, so you'll need to determine the time zone yourself. The `dateutil.tz.gettz` function will return the local time zone of your computer if you pass no argument to it, as shown in Example 4.4

Example 4.4 Obtaining your local time zone

```
1 >>> from dateutil import tz
2 >>> import datetime
3 >>> now = datetime.datetime.now()
4 >>> localzone = tz.gettz()
5 >>> localzone
6 tzfile('/etc/localtime')
7 >>> localzone.tzname(datetime.datetime(2017, 10, 19))
8 'CEST'
9 >>> localzone.tzname(datetime.datetime(2017, 11, 19))
10 'CET'
```

As you can see, we pass two dates to `localzone.tzname(datetime.datetime())` separately, and `dateutil` is able to tell us that one is in Central European Summer Time (CEST) and the other is in Central European Time (no summer). If you pass in your current date, you'll get your own current time zone.

You can use objects from the `dateutil` library in `tzinfo` classes without having to bother implementing those yourself in your application. This makes it easy to convert unaware `datetime` objects to aware `datetime` objects.

4.4. SERIALIZING DATETIME OBJECTS

A class exists in Python that allows you to implement time zone classes yourself: the `datetime.tzinfo` class is an abstract class that provides a base for implementing classes representing time zones. If you ever want to implement a class to represent a time zone, you need to use this as the parent class and implement three different methods:

- `utcoffset(dt)` which must return an offset from UTC in minutes east of UTC for the time zone
- `dst(dt)` which must return the daylight saving time adjustment in minutes east of UTC for the time zone
- `tzname(dt)` which must return the name of the time zone as a string

These three methods will embed a `tzinfo` object, allowing you to translate any time zone—aware datetime to another time zone . However, as mentioned, since those time zone databases exist, it's impractical to implement those time zone classes oneself.

4.4 Serializing datetime Objects

You'll often find you need to transport a `datetime` object from one point to another, where those different points might not be Python native. The typical case nowadays would be with an HTTP REST API, which must return `datetime` objects serialized to a client. The native Python method named `isoformat` can be used to serialize `datetime` objects for non-Python native points, as shown in Example 4.5.

Example 4.5 Serializing a datetime object

```
1 >>> import datetime
2 >>> from dateutil import tz
3 >>> def utcnow():
```

4.4. SERIALIZING DATETIME OBJECTS

```
4     return datetime.datetime.now(tz=tz.tzutc())
5 >>> utcnow()
6 datetime.datetime(2015, 6, 15, 14, 45, 19, 182703, tzinfo=tzutc())
7 >>> utcnow().isoformat()
8 '2015-06-15T14:45:21.982600+00:00'
```

We define a new function called `utcnow` and tell it explicitly to return an object with the UTC time zone. As you can see, the object returned now contains time zone information. We then format the string using the ISO format, ensuring the timestamp also contains some time zone information.

You can see I've used the method `isoformat()` to format the output. I recommend that you always format your datetime input and output strings using ISO 8601, with the method `datetime.datetime.isoformat()`, to return timestamps formatted in a readable way that includes the time zone information.

Your ISO 8601-formatted strings can then be converted to native `datetime.datetime` objects. For this, you can rely on the `iso8601` module to return timestamps with correct time zone information. This module offers only one function, `parse_date`, which does all the hard work of parsing the string and determining the timestamp and time zone values. The fact that the returned timestamps are time zone aware makes it possible to compare them to any other time zone-aware timestamp. The `iso8601` module is not provided as a built-in module in Python, so you need to install it using `pip install iso8601`. Example 4.6 shows how to parse a timestamp using `iso8601`.

Example 4.6 Using the `iso8601` module to parse a ISO 8601 formatted timestamp

```
1 >>> import iso8601
2 >>> import datetime
3 >>> now = datetime.datetime.utcnow()
4 >>> str(now)
5 '2017-06-19 09:42:00.764337'
6 >>> parsed = iso8601.parse_date(now.isoformat())
7 >>> parsed
```

```

8 datetime.datetime(2017, 6, 19, 9, 42, 0, 764337, tzinfo=<iso8601.Utc>)
9 >>> parsed < datetime.datetime.now(tz=pytz.utc)
10 True

```

In Example 4.6, the `iso8601` module is used to construct a `datetime` object from a string. By calling `iso8601.parse_date` on a string containing an ISO 8601 formatted timestamp, the library is able to return a `datetime` object. Since that string does not contain any time zone information, the `iso8601` module assumes that the time zone is UTC. If a time zone contains correct time zone information, the `iso8601` module returns correctly.

Using time zone-aware `datetime` objects and using ISO 8601 as the format for their string representation is a perfect solution for most problems around time zone, making sure no mistakes are made and building great interoperability between your application and the outside world.

4.5 Solving Ambiguous Times

There are certain cases where the time of the day can be ambiguous; for example during the daylight saving time transition when the same “wall clock” time occurs twice a day. The `dateutil` library provides us with the `is_ambiguous` method to distinguish such timestamps. To show this in action, we’ll create an ambiguous timestamp in Example 4.7.

Example 4.7 A confusing timestamp, occurring during the daylight saving time crossover

```

1 >>> import dateutil.tz
2 >>> localtime = dateutil.tz.gettz("Europe/Paris")
3 >>> confusing = datetime.datetime(2017, 10, 29, 2, 30)
4 >>> localtime.is_ambiguous(confusing)
5 True

```

4.6. SUMMARY

On the night of October 30, 2017, Paris switched from summer to winter time. The city switches at 3:00 AM, when the time goes back to 2:00 AM. If we then try to use a timestamp at 2:30 on that date, there is no way for this object to be sure which side of the time zone it is, after or before the daylight saving time change.

However, from Python 3.6, it is impossible to specify which side of the fold a timestamp is using the fold attribute, added to `datetime` objects with the introduction of [PEP 495 \(Local Time Disambiguation\)](#). This attribute indicates on which side of the fold the `datetime` is, as demonstrated in [Example 4.8](#)

Example 4.8 Disambiguating the ambiguous timestamp

```
1 >>> import dateutil.tz
2 >>> import datetime
3 >>> localtime = dateutil.tz.gettz("Europe/Paris")
4 >>> utc = dateutil.tz.tzutc()
5 >>> confusing = datetime.datetime(2017, 10, 29, 2, 30, tzinfo=localtz)
6 >>> confusing.replace(fold=0).astime zone(utc)
7 datetime.datetime(2017, 10, 29, 0, 30, tzinfo=tzutc())
8 >>> confusing.replace(fold=1).astime zone(utc)
9 datetime.datetime(2017, 10, 29, 1, 30, tzinfo=tzutc())
```

You'll need to use this in only very rare cases, since ambiguous timestamps occur only in a small window. Sticking to UTC is a great workaround to keep life simple and avoid running into time zone issues. However, it is good to know that the fold attribute exists and that `dateutil` is able to help in such cases.

4.6 Summary

In this chapter, we have seen how crucial it is to carry this time zone information in timestamps. The built-in `datetime` module is not complete in this regard, but the `dateutil` module is a great complement: it allows us to get `tzinfo`-compatible objects

that are ready to be used. The `dateutil` module also helps us solve subtle issues such as daylight saving time ambiguity.

The ISO 8601 standard format is an excellent choice for serializing and unserializing timestamps because it is readily available in Python and compatible with any other programming language.

CHAPTER 5



Distributing Your Software

It's safe to say that at some point, you will want to distribute your software. As tempted as you might be to just zip your code and upload it to the internet, Python provides tools to make it much easier for your end users to get your software to work. You should already be familiar with using `setup.py` to install Python applications and libraries, but you have probably never delved into how it works behind the scenes, or how to make a `setup.py` of your own.

In this chapter, we'll look into the history and workings of `setup.py` and learn how to create your own for your custom purposes. We'll also take a look at some of the less well-known capabilities of the package installation tool `pip` and how to make your software downloadable via `pip`. Finally, we'll see how to use Python's entry points to make functions easy to find between programs. With these skills, your published software will be accessible for end users.

5.1 A Bit of History

The *distutils* library, originally developed by software developer Greg Ward, has been part of the standard Python library since 1998. Ward sought to create an easy way

for developers to automate the installation process for their end users. The `setup.py` file is the standard Python script provided by packages for their installation. Packages can use *distutils* to get themselves installed themselves, as shown in Example 5.1.

Example 5.1 `setup.py` using *distutils*

```
1  #!/usr/bin/python
2  from distutils.core import setup
3
4  setup(name="rebuildd",
5        description="Debian packages rebuild tool",
6        author="Julien Danjou",
7        author_email="acid@debian.org",
8        url="http://julien.danjou.info/software/rebuildd.html",
9        packages=['rebuildd'])
```

With the `setup.py` file as the root of a project, all users have to do to build or install your software is run that file with the appropriate command as its argument. Even if your distribution includes C modules in addition to native Python ones, *distutils* can handle those automatically as well.

Development of *distutils* was abandoned in 2000; since then, other developers have picked up where it left off, building on its basis. One of the most notable successors is the packaging library known as *setuptools*, which offers more frequent updates and advanced features, such as automatic dependency handling, the Egg distribution format, and the `easy_install` command. Since *distutils* was still the accepted means of packaging software included with the Python Standard Library at the time of development, *setuptools* also provided a degree of backwards compatibility with it. Example 5.2 shows how you'd use *setuptools* to build the same installation package as in Example 5.1.

Example 5.2 `setup.py` using *setuptools*

```
1  #!/usr/bin/env python
2  import setuptools
```

5.1. A BIT OF HISTORY

```
3
4 setuptools.setup(
5     name="pymunincli",
6     version="0.2",
7     author="Julien Danjou",
8     author_email="julien@danjou.info",
9     description="munin client library",
10    license="GPL",
11    url="http://julien.danjou.info/software/pymunincli/",
12    packages=['munin'],
13    classifiers=[
14        "Development Status :: 2 - Pre-Alpha",
15        "Intended Audience :: Developers",
16        "Intended Audience :: Information Technology",
17        "License :: OSI Approved :: GNU General Public License (GPL)",
18        "Operating System :: OS Independent",
19        "Programming Language :: Python"
20    ],
21 )
```

Eventually, development on *setuptools* slowed down too, but it wasn't long before another group of developers forked it to create a new library called *distribute*, which offered several advantages over *setuptools*, including fewer bugs and Python 3 support.

All the best stories have a twist ending, though: in March 2013, the teams behind *setuptools* and *distribute* **decided to merge their code bases** under the aegis of the original *setuptools* project. So *distribute* is now deprecated, and *setuptools* is once more the canonical way to handle advanced Python installations.

While all this was happening, another project, known as *distutils2*, was developed with the intention of replacing *distutils* in the Python Standard Library outright. One of its most notable differences from both *distutils* and *setuptools* was that it stored package metadata in a plain text file, `setup.cfg`, which was easier both for developers to write and for external tools to read. However, *distutils2* retained some of the failings

5.1. A BIT OF HISTORY

of *distutils*, such as its obtuse command-based design, and lacked support for things such as entry points and native script execution on Windows – both features provided by *setuptools*. For these and other reasons, plans to include *distutils2*, renamed as *packaging*, in the Python 3.3 Standard Library fell through, and the project was abandoned in 2012.

There is still, however, a chance for *packaging* to rise from the ashes through *distlib*, an up-and-coming effort to replace *distutils*. Before release, it was rumored that the *distlib* package would become part of the Standard Library in Python 3.4, but that never came to be. Including the best features from *packaging*, *distlib* implements the basic groundwork described in the packaging-related PEPs.

So, to recap:

- *distutils* is part of the Python standard library and can handle simple package installations.
- *setuptools*, the standard for advanced package installations, was at first deprecated but is now back in active development and the de facto standard.
- *distribute* has been merged back into *setuptools* as of version 0.7; *distutils2* (a.k.a. *packaging*) has been abandoned.
- *distlib* might replace *distutils* in the future.

There are other packaging libraries out there, though these are the five you'll encounter the most. Be careful when looking up information about these libraries on the internet: plenty of documentation is outdated due to the complicated history outlined above. The **official documentation** is, at least, up-to-date, however.

In short, *setuptools* is the distribution library to use for the time being, but keep an eye out for *distlib* in the future.

5.2 Packaging with setup.cfg

You've probably already written some package and tried to write a `setup.py` at some point, either by copying one from some other project or by skimming through the documentation and building it yourself. Building a `setup.py` is not an intuitive task. Choosing the right tool to use is just the first challenge. In this section, I want to introduce you to one of `_setuptools_`' recent improvements: the `setup.cfg` file support.

This is what a `setup.py` using a `setup.cfg` file looks like:

```
1 import setuptools
2
3 setuptools.setup()
```

Two lines of code – it is that simple. As promised, the actual metadata the setup requires is stored in `setup.cfg` as in [Example 5.3](#).

Example 5.3 A setup.cfg file

```
1 [metadata]
2 name = foobar
3 author = Dave Null
4 author-email = foobar@example.org
5 license = MIT
6 long_description = file: README.rst
7 url = http://pypi.python.org/pypi/foobar
8 requires-python = >=2.6
9 classifiers =
10     Development Status :: 4 - Beta
11     Environment :: Console
12     Intended Audience :: Developers
13     Intended Audience :: Information Technology
14     License :: OSI Approved :: Apache Software License
15     Operating System :: OS Independent
16     Programming Language :: Python
```

As you can see, `setup.cfg` uses a format that's really easy to write and read, directly inspired by *distutils2*. Many other tools, such as *Sphinx* or *Wheel*, also read configuration from this `setup.cfg` file – that alone is a good argument to start using it.

In Example 5.3, the description of the project is read from the `README.rst` file. It's good practice to always have a `README` file – preferably in the `RST` format – to make it easy to quickly understand what the project is about. With just these basic `setup.py` and `setup.cfg` files, your package is ready to be published and used by other developers and applications. The [setuptools documentation](#) provides more details if needed, for example, if you have some extra steps in your installation process or want to include extra files.

Another packaging tool to know about is *pbr*, short for *Python Build Reasonableness*. The project was started in *OpenStack* as an extension of *setuptools* to facilitate installation and deployment of packages. The *pbr* packaging tool, used alongside *setuptools*, implements features that are absent from *setuptools*, including these:

- Automatic generation of Sphinx documentation
- Automatic generation of `AUTHORS` and `ChangeLog` files based on *git* history
- Automatic creation of file lists for *git*.
- Version management based on *git* tags using semantic versioning.

And all this with little to no effort on your part. To use *pbr*, you just need to enable, it, as shown in Example 5.4.

Example 5.4 `setup.py` using *pbr*

```

1 import setuptools
2
3 setuptools.setup(setup_requires=['pbr'], pbr=True)
```

The `setup_requires` parameter indicates to *setuptools* that *pbr* must be installed prior to using *setuptools*. The `pbr=True` argument makes sure that the *pbr* extension for *setuptools* is loaded and called.

Once enabled, the `python setup.py` command is enhanced with the *pbr* features. Calling `python setup.py --version` will, for example, return the version number of the project based on existing *git* tags. Running `python setup.py sdist` would create a source tarball with automatically generated `ChangeLog` and `AUTHORS` files.

5.3 The *Wheel* Format Distribution Standard

For most of Python's existence, there's been no official standard distribution format. While different distribution tools still generally use some common archive format – even the *Egg* format introduced by *setuptools* is just a zip file with a different extension – their metadata and package structures are incompatible with each other. This problem was compounded when an official installation standard was finally defined in [PEP 376](#), that was also incompatible with existing formats.

To solve these problems, [PEP 427](#) was written to define a new standard for Python distribution packages called *Wheel*. The reference implementation of this format is available as a tool, also called [wheel1](#).

Wheel is supported by [pip](#) starting with version 1.4. If you are using *setuptools* and have the *wheel* package installed, it automatically integrates itself as a *setuptools* command named `bdist_wheel`. If you don't have *Wheel* installed, you can install it using the command `pip install wheel`. Calling `setup.py bdist_wheel` shows some of the output when calling `bdist_wheel`, abridged for print.

Calling `setup.py bdist_wheel`

```
1 $ python setup.py bdist_wheel
2 running bdist_wheel
3 running build
```

5.3. THE *WHEEL* FORMAT DISTRIBUTION STANDARD

```
4 running build_py
5 creating build/lib
6 creating build/lib/daiquiri
7 creating build/lib/daiquiri/tests
8 copying daiquiri/tests/__init__.py -> build/lib/daiquiri/tests
9 copying daiquiri/tests/test_daiquiri.py -> build/lib/daiquiri/tests
10 [...]
11 running egg_info
12 writing requirements to daiquiri.egg-info/requires.txt
13 writing daiquiri.egg-info/PKG-INFO
14 writing top-level names to daiquiri.egg-info/top_level.txt
15 writing dependency_links to daiquiri.egg-info/dependency_links.txt
16 writing pbr to daiquiri.egg-info/pbr.json
17 writing manifest file 'daiquiri.egg-info/SOURCES.txt'
18 installing to build/bdist.macosx-10.12-x86_64/wheel
19 running install
20 running install_lib
21 [...]
22 running install_egg_info
23 Copying daiquiri.egg-info to build/bdist.macosx-10.12-x86_64/wheel/ ↵
    daiquiri-1.3.0-py2.7.egg-info
24 running install_scripts
25 creating build/bdist.macosx-10.12-x86_64/wheel/daiquiri-1.3.0.dist- ↵
    info/WHEEL
26 creating '/Users/jd/Source/daiquiri/dist/daiquiri-1.3.0-py2.py3-none- ↵
    any.whl' and adding '.' to it
27 adding 'daiquiri/__init__.py'
28 [...]
```

The `bdist_wheel` command creates a `.whl` file in the `dist` directory. As with the *Egg* format, a *Wheel* archive is just a zip file with a different extension. However, *Wheel* archives do not require installation – you can load and run your code just by adding a slash followed by the name of your module:

```
1 $ python wheel-0.21.0-py2.py3-none-any.whl/wheel -h
2 usage: wheel [-h]
3
```


5.3. THE *WHEEL* FORMAT DISTRIBUTION STANDARD

```
4         {keygen,sign,unsign,verify,unpack,install,install-scripts, ↵  
          convert,help}  
5         ...  
6  
7 positional arguments:  
8 [...]
```

You might be surprised to learn this is not a feature introduced by the *Wheel* format itself. Python can also run regular zip files as well, just like with Java's `.jar` files:

```
1 python foobar.zip
```

is equivalent to:

```
1 PYTHONPATH=foobar.zip python -m __main__
```

In other words, the `__main__` module for your program will automatically be imported from `__main__.py`. It is also possible to import `__main__` from a module you specify by appending a slash followed by its name, just like with *Wheel*:

```
1 python foobar.zip/mymod
```

This is equivalent to:

```
1 PYTHONPATH=foobar.zip python -m mymod.__main__
```

One of the advantages of *Wheel* is that its naming conventions allow you to specify whether your distribution is intended for a specific architecture and Python implementation (CPython, PyPy, Jython, etc.). This is particularly useful if you need to distribute modules written in C.

So by default, *Wheel* packages are tied to the major version of Python that you used to build them. When called with `python2 setup.py bdist_wheel`, the pattern of a wheel file name will be something like `library-version-py2-none-any.whl`.

If your code is compatible with all major Python versions, that is Python 2 and Python 3, you can build a *universal* wheel:

```
1 python setup.py bdist_wheel --universal
```

The resulting filename will be different and contains both Python major versions, e.g., `library-version-py2.py3-none-any.whl`. Building a universal wheel avoids having 2 different wheels when only one would cover both Python major versions.

**Tip**

If you do not want to pass the `--universal` flag each time you are building a wheel, you can add this to your `setup.cfg` file.

```
1 [wheel]
2 universal=1
```

If the wheel you build contains binary programs or libraries (like a Python extension written in C), the binary wheel might not be as portable as you can imagine. It will work by default on some platforms, such as Darwin (macOS) or Microsoft Windows, but it might not work on all Linux distributions. The [PEP 513](#) targets this Linux problem by defining a new platform tag named `manylinux1` and a minimal set of libraries that are guaranteed to be available on that platform.

Wheels is a great format for the distribution ready-to-install libraries and applications, therefore it is encouraged to build and upload them to PyPI as well.

5.4 Sharing Your Work with the World

Once you have a proper `setup.py` file, it is easy to build a source tarball that can be distributed. The `sdist` *setuptools* command does just that, as demonstrated in [Example 5.5](#).

Example 5.5 Using `setup.py sdist`

```

1 $ python setup.py sdist
2 running sdist
3
4 [pbr] Generating AUTHORS
5 running egg_info
6 writing requirements to ceilometer.egg-info/requirements.txt
7 writing ceilometer.egg-info/PKG-INFO
8 writing top-level names to ceilometer.egg-info/top_level.txt
9 writing dependency_links to ceilometer.egg-info/dependency_links.txt
10 writing entry points to ceilometer.egg-info/entry_points.txt
11 [pbr] Processing SOURCES.txt
12 [pbr] In git context, generating filelist from git
13 warning: no previously-included files matching '*.pyc' found anywhere ↵
    in distribution
14 writing manifest file 'ceilometer.egg-info/SOURCES.txt'
15 running check
16 copying setup.cfg -> ceilometer-2014.1.a6.g772e1a7
17 Writing ceilometer-2014.1.a6.g772e1a7/setup.cfg
18
19 [...]
20
21 Creating tar archive
22 removing 'ceilometer-2014.1.a6.g772e1a7' (and everything under it)

```

The `sdist` command creates a tarball under the `dist` directory of the source tree. The tarball contains all the Python modules that are part of the source tree. As seen in Section 5.3, it is also possible to build *Wheel* archives using the `bdist_wheel` command, which are a bit faster to install as they are already in the right format for that.

The final step to make that code accessible is to make things easy on your end-users by exporting your package where users can install it via *pip*. That means publishing your project to **PyPI**.

5.4. SHARING YOUR WORK WITH THE WORLD

Since you might make mistakes if this is your first time, it pays to test out the publishing process in a safe sandbox rather than on the production server. You can use the **PyPI staging server** for this purpose: it replicates all the functionality of the main index, but it is used solely for testing purposes.

The first step is to register your project on the test server. Start by opening your `~/.pypirc` file and adding these lines:

```
1 [distutils]
2 index-servers =
3     testpypi
4
5 [testpypi]
6 username = <your username>
7 password = <your password>
8 repository = https://testpypi.python.org/pypi
```

Save the file, and now you can register your project in the index:

```
1 $ python setup.py register -r testpypi
2 running register
3 running egg_info
4 writing requirements to ceilometer.egg-info/requires.txt
5 writing ceilometer.egg-info/PKG-INFO
6 writing top-level names to ceilometer.egg-info/top_level.txt
7 writing dependency_links to ceilometer.egg-info/dependency_links.txt
8 writing entry points to ceilometer.egg-info/entry_points.txt
9 [pbr] Reusing existing SOURCES.txt
10 running check
11 Registering ceilometer to https://testpypi.python.org/pypi
12 Server response (200): OK
```

This connects to the test PyPI server instance and creates a new entry. Do not forget to use the `-r` option; otherwise, the real production PyPI instance would be used!

Obviously, if a project with the same name is already registered there, the process will fail. Retry with a new name, and once you get your program registered and

5.4. SHARING YOUR WORK WITH THE WORLD

receive the OK response, you can upload a source distribution tarball:

```
1 % python setup.py sdist upload -r testpypi
2 running sdist
3 [pbr] Writing ChangeLog
4 [pbr] Generating AUTHORS
5 running egg_info
6 writing requirements to ceilometer.egg-info/requires.txt
7 writing ceilometer.egg-info/PKG-INFO
8 writing top-level names to ceilometer.egg-info/top_level.txt
9 writing dependency_links to ceilometer.egg-info/dependency_links.txt
10 writing entry points to ceilometer.egg-info/entry_points.txt
11 [pbr] Processing SOURCES.txt
12 [pbr] In git context, generating filelist from git
13 warning: no previously-included files matching '*.pyc' found anywhere ↵
    in distribution
14 writing manifest file 'ceilometer.egg-info/SOURCES.txt'
15 running check
16 creating ceilometer-2014.1.a6.g772e1a7
17
18 [...]
19
20 copying setup.cfg -> ceilometer-2014.1.a6.g772e1a7
21 Writing ceilometer-2014.1.a6.g772e1a7/setup.cfg
22 Creating tar archive
23 removing 'ceilometer-2014.1.a6.g772e1a7' (and everything under it)
24 running upload
25 Submitting dist/ceilometer-2014.1.a6.g772e1a7.tar.gz to https:// ↵
    testpypi.python.org/pypi
26 Server response (200): OK
```

Alternatively, you could upload a *Wheel* archive:

```
1 $ python setup.py bdist_wheel upload -r testpypi
2 running bdist_wheel
3 running build
4 running build_py
5 running egg_info
6 writing requirements to ceilometer.egg-info/requires.txt
```

5.4. SHARING YOUR WORK WITH THE WORLD

```
7 writing ceilometer.egg-info/PKG-INFO
8 writing top-level names to ceilometer.egg-info/top_level.txt
9 writing dependency_links to ceilometer.egg-info/dependency_links.txt
10 writing entry points to ceilometer.egg-info/entry_points.txt
11 [pbr] Reusing existing SOURCES.txt
12 installing to build/bdist.linux-x86_64/wheel
13 running install
14 running install_lib
15 creating build/bdist.linux-x86_64/wheel
16
17 [...]
18
19 creating build/bdist.linux-x86_64/wheel/ceilometer-2014.1.a6.g772e1a7. ←
    dist-info/WHEEL
20 running upload
21 Submitting /home/jd/Source/ceilometer/dist/ceilometer-2014.1.a6. ←
    g772e1a7-py27-none-any.whl to https://testpypi.python.org/pypi
22 Server response (200): OK
```

Once those operations are finished, you and other users can search for the uploaded packages on the [PyPi staging server](#), and even install those packages using *pip*, by specifying the test server using the `-i` option:

```
1 $ pip install -i https://testpypi.python.org/pypi ceilometer
```

If everything checks out, you can continue to the next step: uploading your project to the main PyPI server. Just add your credentials and the details for the server to your `~/.pypirc` file, like so:

```
1 [distutils]
2 index-servers =
3     pypi
4     testpypi
5
6 [pypi]
7 username = <your username>
8 password = <your password>
```

```

9 [testpypi]
10 repository = https://testpypi.python.org/pypi
11 username = <your username>
12 password = <your password>
13

```

Now if you run *register* and upload with the `-r pypi` switch, your package should be uploaded to PyPI proper.



Tip

PyPI can keep several versions of your software in its index, allowing you to install specific and older versions, if ever needed to. Just pass the version number to the `pip install` command, for example, `pip install foobar==1.0.2`.

This process is straightforward to use and allows for any number of uploads. You can release your software as often that you want, and your users can install and update as often as they need.

5.5 Entry Points

You may have already used *setuptools* entry points without knowing anything about them.

Software distributed using *setuptools* includes important metadata describing things such as its required dependencies and – more relevantly to this topic – a list of “entry points”. Entry points are methods by which, other Python programs can discover the dynamical features that a package provides.

The following is an example of providing example shows how to provide an entry point named `rebuild` in the `console_scripts` entry point group:

```

1  #!/usr/bin/python
2  from distutils.core import setup
3
4  setup(name="rebuild",
5      description="Debian packages rebuild tool",
6      author="Julien Danjou",
7      author_email="acid@debian.org",
8      url="http://julien.danjou.info/software/rebuild.html",
9      entry_points={
10         'console_scripts': [
11             'rebuild = rebuild:main',
12         ],
13     },
14     packages=['rebuild'])

```

Any Python package can register entry points. Entry points are organized in groups: each group is made of a list of key and value pairs. Those pairs use the format `path.to.module:variable_name`. In the example above, the key is `rebuild`, and the value is `rebuild:main`.

The list of entry points can be manipulated using various tools, from *setuptools* to *epi*, as I'll show here. In the following sections, we discuss how we can entry points to add extensibility to our software.

5.5.1 Visualising Entry Points

The easiest way to visualize the entry points available in a package is to use a package called *entry_point_inspector*. You can install it by running `pip install entry-point-inspector`.

When installed, it provides the command `epi` that you can run from your terminal to interactively discover the entry points provided by installed packages. Example 5.6 shows an example of running `epi group list` on my system.

Example 5.6 Getting list of entry points groups

```
1 $ epi group list
2 +-----+
3 | Name                |
4 +-----+
5 | console_scripts      |
6 | distutils.commands   |
7 | distutils.setup_keywords |
8 | egg_info.writers     |
9 | epi.commands         |
10 | flake8.extension     |
11 | setuptools.file_finders |
12 | setuptools.installation |
13 +-----+
```

The output from `epi group list` here in Example 5.6 shows the different packages on a system that provide entry points. Each item in this table is the name of an entry point group. Note that this list includes `console_scripts`, which we’ll discuss shortly. We can use the `epi` command with the `show` command to show details of a particular entry point group with the `show` command, as in Example 5.7.

Example 5.7 Showing details of an entry point group

```
1 $ epi group show console_scripts
2 +-----+-----+-----+-----+-----+
3 | Name      | Module   | Member | Distribution | Error |
4 +-----+-----+-----+-----+-----+
5 | coverage | coverage | main   | coverage 3.4 |      |
6 +-----+-----+-----+-----+-----+
```

Here we can see that in the group `console_scripts` , an entry point named `coverage` refers to the member `main` of the module `coverage`. This entry point in particular, provided by the package `coverage 3.4`, indicates which Python function to call when the command line script `coverage` is executed. Here, the function `cove`

`rage.main` is to be called.

The `epi` tool used here is just a thin layer on top of the complete Python library `pkg_resources`. This module allows us to discover entry points for any Python library or program. Entry points are valuable for various things, including console scripts and dynamic code discovery, as we are going to see in the next few sections.

5.5.2 Using Console Scripts

When writing a Python application, you almost always have to provide a launchable program – a Python script that the end user can run. This program needs to be installed inside a directory somewhere in the system path.

Most projects have a launchable program similar to this:

```
1  #!/usr/bin/python
2  import sys
3  import mysoftware
4
5  mysoftware.SomeClass(sys.argv).run()
```

This kind of script is a best-case scenario: many projects have a much longer script installed in the system path. However, such scripts pose some major issues:

- There's no way they can know where the Python interpreter is or which version it is.
- They leak binary code that can't be imported by software or unit tests.
- There's no easy way to define where to install them.
- It's not obvious how to install this in a portable way (for example, on both Unix and Windows).

5.5. ENTRY POINTS

Helping us circumvent these problems, *setuptools* offers the `console_scripts` feature. This entry point can be used to make *setuptools* install a tiny program in the system path that then calls a specific function in one of your modules. With *setuptools*, you can specify a function call to use to start your program by setting up a key/value pair in the `console_scripts` entry point group: the key is the script name that will be installed, and the value is the Python path to your function (something like `my_module.main`).

Let's imagine a *foobar* program that consists of a client and a server. Each part is written in its module – *foobar.client* and *foobar.server*, respectively:

Example 5.8 foobar/client.py

```
1 def main():
2     print("Client started")
```

Example 5.9 foobar/server.py

```
1 def main():
2     print("Server started")
```

Of course, this program doesn't do much of anything – our client and server don't even talk to each other. For our example, though, they just need to print a message letting us know they have started successfully.

We can now write the following `setup.py` file in the root directory with entry points defined:

setup.py

```
1 from setuptools import setup
2
3 setup(
4     name="foobar",
5     version="1",
6     description="Foo!",
```

5.5. ENTRY POINTS

```
7     author="Julien Danjou",
8     author_email="julien@danjou.info",
9     packages=["foobar"],
10    entry_points={
11        "console_scripts": [
12            "foobard = foobar.server:main",
13            "foobar = foobar.client:main",
14        ],
15    },
16 )
```

We define entry points using the format `module.submodule:function`. You can see here that we've defined an entry point for both `client` and `server`.

When `python setup.py install` is ran, *setuptools* will create a script that will look like the one in Example 5.10.

Example 5.10 A console script generated by *setuptools*

```
1  #!/usr/bin/python
2  # EASY-INSTALL-ENTRY-SCRIPT: 'foobar==1','console_scripts','foobar'
3  __requires__ = 'foobar==1'
4  import sys
5  from pkg_resources import load_entry_point
6
7  if __name__ == '__main__':
8      sys.exit(
9          load_entry_point('foobar==1', 'console_scripts', 'foobar')()
10     )
```

This code scans the entry points of the `foobar` package and retrieves the `foobar` key from the `console_scripts` group, which is used to locate and run the corresponding function. The return value of the `load_entry_point` will then be a reference to the function `foobar.client.main`, that will be called without any argument and whose return value will be used as an exit code. Notice that this code uses `pkg_resources` to discover and load entry point files from within your Python programs.

**Tip**

If you're using *pbr* on top of *setuptools*, the generated script is simpler (and therefore faster) than the default one built by *setuptools* as it will call the function you wrote in the entry point without having to search the entry point list dynamically at runtime.

Using console scripts is a great technique that remove the burden of writing portable scripts, while ensuring that your code stays in your Python package and can be imported (and tested) by other programs.

5.5.3 Using Plugins and Drivers

Entry points make it easy to discover and dynamically load code deployed by other packages, but this is not the only use of entry points. Any application can propose and register entry points and groups and then use them as it wishes.

In this section, we're going to create a *cron*-style daemon that will allow any Python program to register a command to be run once every few seconds by registering an entry point in the group `pytimed`. The attribute pointed to by this entry point should be an object that returns `number_of_seconds`, callable.

Here's our implementation using `pkg_resources` to discover entry points, in a program I've named `pytimed.py`.

pytimed.py

```
1 import pkg_resources
2 import time
3
4 def main():
5     seconds_passed = 0
6     while True:
7         for entry_point in pkg_resources.iter_entry_points('pytimed'):
```

5.5. ENTRY POINTS

```
8         try:
9             seconds, callable = entry_point.load()()
10        except:
11            # Ignore failure
12            pass
13        else:
14            if seconds_passed % seconds == 0:
15                callable()
16        time.sleep(1)
17        seconds_passed += 1
```

This program consists of an infinite loop that iterates over each entry point of the `pytimed` group. Each entry point is loaded using the `load()` method. The returned method is then called, and it needs to return both a number of seconds to wait before calling the callable, and the aforementioned callable.

The program in `pytimed.py` is a very simplistic and naive implementation, but it is sufficient for our example. Now we can write another Python program, named `hello.py`, that needs one of its functions called on a periodic basis:

hello.py

```
1 def print_hello():
2     print("Hello, world!")
3
4 def say_hello():
5     return 2, print_hello
```

Once we have that function defined, we register the function using the appropriate entry points in `setup.py`:

setup.py

```
1 from setuptools import setup
2
3 setup(
4     name="hello",
5     version="1",
```

5.5. ENTRY POINTS

```
6 packages=["hello"],
7 entry_points={
8     "pytimed": [
9         "hello = hello:say_hello",
10    ],
11 },)
```

The `setup.py` script registers an entry point in the group `pytimed` with the key `hello` and the value pointing to the function `hello.say_hello`. Once that package is installed using that `setup.py` – for example, using `pip install` – the `pytimed` script can detect the newly added entry point.

At startup, `pytimed` will scan the group `pytimed` and find the key `hello`. It will then call the `hello.say_hello` function, getting two values: the number of seconds to wait between each call, 2 seconds in this case, and the function to call, `print_hello` in this case. By running the program, as we do in Example 5.11, you can see "Hello, world!" printed on the screen every two 2 seconds.

Example 5.11 Running `pytimed`

```
1 % python3
2 Python 3.3.2+ (default, Aug 4 2013, 15:50:24)
3 [GCC 4.8.1] on linux
4 Type "help", "copyright", "credits" or "license" for more information.
5 >>> import pytimed
6 >>> pytimed.main()
7 Hello, world!
8 Hello, world!
9 Hello, world!
```

The possibilities this mechanism offers are immense: it allows you to build driver systems, hook systems, and extensions easily and generically. Implementing this mechanism by hand in every program you make would be tedious, but fortunately, there's a Python library that can take care of the boring parts for us.

5.5. ENTRY POINTS

The **stevedore** library provides support for dynamic plugins based on the same mechanism demonstrated in our previous examples. The use case in this example is already simplistic, but we can still simplify it a bit further using **stevedore**:

pytimed_stevedore.py

```
1 from stevedore.extension import ExtensionManager
2 import time
3
4 def main():
5     seconds_passed = 0
6     extensions = ExtensionManager('pytimed', invoke_on_load=True)
7     while True:
8         for extension in extensions:
9             try:
10                 seconds, callable = extension.obj
11             except:
12                 # Ignore failure
13                 pass
14             else:
15                 if seconds_passed % seconds == 0:
16                     callable()
17             time.sleep(1)
18             seconds_passed += 1
```

The `ExtensionManager` class of `stevedore` provides a simple way to load all extensions of an entry point group. The name is passed as a first argument. The argument `invoke_on_load=True` makes sure that each function of the group is called once discovered. This makes the results accessible directly from the `obj` attribute of the extension.

If you look through the *stevedore* documentation, you will see that `ExtensionManager` has a variety of subclasses that can handle different situations, such as loading specific extensions based on their names or the result of a function. All of those are commonly used models that you can apply easily to your program in order to implement those patterns directly.

5.5. ENTRY POINTS

For example, we might want to load and run only one extension, rather than all of them, from our entry point group. Leveraging the `stevedore.driver.DriverManager` class allows us to do that, as Example 5.12 shows.

Example 5.12 Using *stevedore* to run a single extension from an entry point

```
1 from stevedore.driver import DriverManager
2 import time
3
4 def main(name):
5     seconds_passed = 0
6     seconds, callable = DriverManager('pytimed', name, invoke_on_load= ↵
7         True).driver
8     while True:
9         if seconds_passed % seconds == 0:
10             callable()
11             time.sleep(1)
12             seconds_passed += 1
13
14 main("hello")
```

In this case, only one extension is loaded and selected by its name. This allows to quickly build a driver system in which only one extension is loaded and used by a program.

5.5.4 Summary

The packaging ecosystem in Python has a bumpy history, however things are now getting settled. *setuptools* provide a complete solution to packaging, not only to transport your code in different formats and upload it to PyPI, but also to handle connection with other software and libraries via entry points.

5.6 Interview with Nick Coghlan

Nick Coghlan is a Python core developer working at Red Hat. He has written several PEP proposals, including [PEP 426](#) (*Metadata for Python Software Packages 2.0*) for which he is acting as *BDFL*¹ delegate.

“



The number of packaging solutions (*distutils*, *setuptools*, *distutils2*, *distlib*, *bento*, *pbr*, etc.) for Python is quite impressive. In your opinion, what are the (possibly historical) reasons for such fragmentation and divergence?

The short answer is that software publication, distribution, and integration is a complex problem with plenty of room for multiple solutions tailored for different use cases. The long answer can be found in the [Python Packaging User Guide](#). In my recent talks on this, I have noted that the problem is mainly one of age with the tools mentioned above being born in a somewhat different era of software distribution.

***setuptools* is the *de facto* standard for Python distributions nowadays. Is there anything you think users should be aware of when using it (or not)?**

¹“Benevolent Dictator For Life,” title given to Guido van Rossum, author of Python

setuptools is entirely reasonable as a build system, especially for pure Python projects, or those with only simple C extensions. It also offers a powerful system for plugin registration and good cross-platform script generation.

While effective, the multi-version support in *pkg_resources* is also a bit quirky and tricky to use properly. Unless there is a compelling reason to have conflicting versions in the same environment, it is much easier to just use *virtualenv* or *zc.buildout*.

PEP 426, which defines a new metadata format for Python packages, is still fairly recent and not yet approved. Is it on good tracks? What motivated it in the first place, how do you think it will tackle the current problems?

PEP 426 originally started as part of the *Wheel* format definition, but Daniel Holth eventually realized that *Wheel* could work with the existing metadata format defined by *setuptools*. PEP 426 is thus a consolidation of the existing *setuptools* metadata with some of the ideas from *distutils2* and other packaging systems (like *RPM* and *npm*), and also addresses some of the frustrations encountered with existing tools (like cleanly separating different kinds of dependencies).

If PEP 426 is accepted, what kinds of tools would you to see built to take advantage of what it offers?

The main gains will be a REST API on PyPI offering full metadata access, as well as (hopefully) the ability to automatically generate distribution policy-compliant packages from upstream metadata.

The *Wheel* format is somewhat recent and not widely used yet, but it seems promising. Is there any reason it is not part of the Standard Library, or are there already plans to include it?

It turns out the Standard Library is not really a suitable place for packaging standards: it evolves too slowly, and an addition to a later version of the Standard Library cannot be used with earlier versions of Python. So, at the Python language summit earlier this year, we tweaked the PEP process to allow *distutils-sig* to manage the full approval cycle for packaging-related PEPs. *python-dev* will only be involved for proposals that involve changing CPython directly (like *pip* bootstrapping).

What kind of future do you envision that would push developers to build and distribute *Wheel* packages?

We still have some tweaks to make before *Wheel* is suitable for use on Linux. However, *pip* is adopting it as an alternative to the *Egg* format, allowing local caching of builds for fast virtual environment creation, and PyPI allows uploads of *Wheel* archives for Windows and macOS.

”

CHAPTER 6



Unit Testing

Many find unit testing to be arduous and time-consuming, and you'll even find people and projects with no testing policy. This chapter assumes that you see the wisdom of unit testing! Writing code that is not tested is fundamentally useless, as there's no way to conclusively prove that it works, but if you need convincing, I suggest you start by reading about the benefits of test-driven development.

The purpose of this chapter is to walk you through the Python tools you can use to construct a comprehensive suite of tests that will make testing simpler and more automated. We'll talk about how you can use tools to enhance your software, making it rock solid and regression-free. We'll cover creating reusable test objects, running tests in parallel, uncovering untested code, and using virtual environments to make sure your tests are clean, as well as some other good-practice methods and ideas.

6.1 The Basics of Testing

Contrary to what you may believe, writing and running unit tests is uncomplicated in Python. The process is not intrusive or disruptive, and unit testing will greatly help

you and other developers in maintaining your software. Here I'll discuss some of the absolute basics of testing that will make things easier for you.

6.1.1 Some Simple Tests

First, you should store tests inside a `tests` submodule of the application or library they apply to. Doing so will allow you to ship the tests as part of your module so that they can be run or reused by anyone—even after your software is installed—without necessarily using the source package. Making the tests a submodule of your main module also prevents them from being installed by mistake in a top-level tests module.

Using a hierarchy in your test tree that mimics the hierarchy of your module tree will make the tests more manageable. This means that the tests covering the code of `mylib/foobar.py` should be stored inside `mylib/tests/test_foobar.py`. Consistent nomenclature will make things simpler when you're looking for the tests related to a particular file. Example 6.1 shows the simplest unit test you can write.

Example 6.1 A really simple test in `test_true.py`

```
1 def test_true():  
2     assert True
```

This will simply assert that the behavior you expect is the behavior of the program. To run this test, you need to load the `test_true.py` file and run the `test_true()` function defined within.

However, writing and running an individual test for each of your test files and functions would clearly be a pain. For small projects with simple usage, the `pytest` package comes to the rescue—once installed via `pip`, `pytest` provides the `pytest` command, which loads every file whose name starts with `test_` and then executes all functions within that start with `test_`.

6.1. THE BASICS OF TESTING

With just the `test_true.py` file in our source tree, running `pytest` will give us the following output:

```
1 $ pytest -v test_true.py
2 ===== test session starts ⌵
   =====
3 platform darwin -- Python 3.6.4, pytest-3.3.2, py-1.5.2, pluggy-0.6.0 ⌵
   -- /usr/ local/opt/python/bin/python3.6
4 cachedir: .cache
5 rootdir: examples, inifile:
6 collected 1 item
7
8 test_true.py::test_true PASSED [100%]
9
10 ===== 1 passed in 0.01 seconds ⌵
   =====
```

The `-v` option tells `pytest` to be verbose and print the name of each test run on one line. If a test fails, the output changes to indicate the failure, accompanied by the whole traceback.

Let's add a failing test this time, as shown in [Example 6.2](#).

Example 6.2 A failing test in `test_true.py`

```
1 def test_false():
2     assert False
```

If we run the test file again, here what happens:

```
1 $ pytest -v test_true.py
2 ===== test session starts ⌵
   =====
3 platform darwin -- Python 3.6.4, pytest-3.3.2, py-1.5.2, pluggy-0.6.0 ⌵
   -- /usr/ local/opt/python/bin/python3.6
4 cachedir: .cache
5 rootdir: examples, inifile:
6 collected 2 items
```

6.1. THE BASICS OF TESTING

```
7
8 test_true.py::test_true PASSED [ 50%]
9 test_true.py::test_false FAILED [100%]
10
11 ===== FAILURES ↵
12 ----- test_false ↵
13
14 def test_false():
15 >     assert False
16 E     assert False
17
18 test_true.py:5: AssertionError
19 ===== 1 failed, 1 passed in 0.07 seconds ↵
20 =====
```

A test fails as soon as an `AssertionError` exception is raised; our `assert` test will raise an `AssertionError` when its argument is evaluated to something false (`False`, `None`, `0`, etc.). If any other exception is raised, the test also errors out.

Simple, isn't it? While simplistic, this approach is used by a lot of small projects and works very well. Those projects require no tools or libraries other than `pytest` and thus can rely on simple `assert` tests.

As you start to write more sophisticated tests, `pytest` will help you understand what's wrong in your failing tests. Imagine the following test:

```
1 def test_key():
2     a = ['a', 'b']
3     b = ['b']
4     assert a == b
```

When `pytest` is run, it gives the following output:

```
1 $ pytest test_true.py
2 ===== test session starts ↵
3 =====
```


6.1. THE BASICS OF TESTING

```
3 platform darwin -- Python 3.6.4, pytest-3.3.2, py-1.5.2, pluggy-0.6.0
4 rootdir: /Users/jd/Source/python-book/examples, inifile:
5 plugins: celery-4.1.0
6 collected 1 item
7
8 test_true.py F [100%]
9
10 ===== FAILURES ↵
11
12
13 ----- test_key ↵
14
15
16 def test_key():
17     a = ['a', 'b']
18     b = ['b']
19     > assert a == b
20 E     AssertionError: assert ['a', 'b'] == ['b']
21 E         At index 0 diff: 'a' != 'b'
22 E         Left contains more items, first extra item: 'b'
23 E         Use -v to get the full diff
24
25 test_true.py:10: AssertionError
26 ===== 1 failed in 0.07 seconds ↵
27
28 =====
```

This tells us that `a` and `b` are different and that this test does not pass. But it also tells us exactly how they are different, making it easy to fix the test or code.

6.1.2 Skipping Tests

If a test cannot be run, you will probably want to skip that test — for example, you may wish to run a test conditionally based on the presence or absence of a particular library. To that end, you can use the `pytest.skip()` function, which will mark the test as skipped and move on to the next one. The `pytest.mark.skip` decorator skips

the decorated test function unconditionally, so you'll use it when a test always needs to be skipped. Example 6.3 shows how to skip a test using these various methods.

Example 6.3 Skipping tests

```

1 import pytest
2
3 try:
4     import mylib
5 except ImportError:
6     mylib = None
7
8
9 @pytest.mark.skip("Do not run this")
10 def test_fail():
11     assert False
12
13
14 @pytest.mark.skipif(mylib is None, reason="mylib is not available")
15 def test_mylib():
16     assert mylib.foobar() == 42
17
18
19 def test_skip_at_runtime():
20     if True:
21         pytest.skip("Finally I don't want to run it")
  
```

When executed, this test file will output the following:

```

1 $ pytest -v examples/test_skip.py
2 ===== test session starts =====
3 platform darwin -- Python 3.6.4, pytest-3.3.2, py-1.5.2, pluggy-0.6.0 ←
4 -- /usr/local/opt/python/bin/python3.6
5 cachedir: .cache
6 rootdir: examples, inifile:
7 collected 3 items
8
9 examples/test_skip.py::test_fail SKIPPED [ 33%]
  
```

```

9  examples/test_skip.py::test_mylib SKIPPED [ 66%]
10 examples/test_skip.py::test_skip_at_runtime SKIPPED [100%]
11
12 ===== 3 skipped in 0.01 seconds =====

```

The output of the test run in Example 6.3 indicates that, in this case, all the tests have been skipped. Having this information reported means you can be sure you didn't accidentally skip a test you expected to run while developing.

6.1.3 Running Particular Tests

When using `pytest`, you often want to run only a particular subset of your tests. You can select which tests you want to run by passing their directory or files as an argument to the `pytest` command line. For example, calling `pytest test_one.py` will only run the `test_one.py` test. `Pytest` also accepts a directory as argument, and in that case, it will recursively scan the directory and run any file that matches the `test_*.py` pattern.

You can also add a filter with the `-k` argument on the command line in order to execute only the test matching a name, as shown in Example 6.4.

Example 6.4 Filtering tests run by name

```

1  $ pytest -v examples/test_skip.py -k test_fail
2  ===== test session starts =====
3  =====
4  platform darwin -- Python 3.6.4, pytest-3.3.2, py-1.5.2, pluggy-0.6.0 ←
   -- /usr/
5  local/opt/python/bin/python3.6
6  cachedir: .cache
7  rootdir: examples, inifile:
8  collected 3 items
9
10 examples/test_skip.py::test_fail SKIPPED [100%]

```

6.1. THE BASICS OF TESTING

```
11  
12 === 2 tests deselected ===  
13 === 1 skipped, 2 deselected in 0.04 seconds ===
```

Names are not always the best way to filter which tests will run. Commonly, a developer would group tests by functionalities or types instead. Pytest provides a dynamic marking system that allows you to mark tests with a key- word that can be used as a filter. To mark tests in this way, use the `-m` option. So if we set up a couple of tests like so...

```
1 import pytest  
2  
3 @pytest.mark.dicctest  
4 def test_something():  
5     a = ['a', 'b']  
6     assert a == a  
7  
8 def test_something_else():  
9     assert False
```

... we can use the `-m` argument with pytest to run only one of those tests:

```
1 $ pytest -v test_mark.py -m dicctest  
2 ===== test session starts =====  
3 platform darwin -- Python 3.6.4, pytest-3.3.2, py-1.5.2, pluggy-0.6.0 ←  
   -- /usr/ local/opt/python/bin/python3.6  
4 cachedir: .cache  
5 rootdir: examples, inifile:  
6 collected 2 items  
7  
8 test_mark.py::test_something PASSED      [100%]  
9  
10 === 1 tests deselected ===  
11 === 1 passed, 1 deselected in 0.01 seconds ===
```

The `-m` marker accepts more complex queries, so we can also run all tests that are *not* marked:

```

1 $ pytest test_mark.py -m 'not dicttest'
2 === test session starts ===
3 platform darwin -- Python 3.6.4, pytest-3.3.2, py-1.5.2, pluggy-0.6.0 ←
   rootdir: examples, inifile:
4 collected 2 items
5 test_mark.py F
6
7 [100%]
8 === FAILURES ===
9 test_something_else
10 22 Chapter 6
11 def test_something_else():
12 >     assert False
13 E     assert False
14 test_mark.py:10: AssertionError
15 === 1 tests deselected ===
16 === 1 failed, 1 deselected in 0.07 seconds ===

```

Here pytest executed every test that was not marked as `dicttest` — in this case, the `test_something_else` test, which failed. The remaining marked test, `test_something`, was not executed and so is listed as deselected.

Pytest accepts more complex expressions—composed of the `or`, `and`, and `not` keywords — allowing you to do more advanced filtering to customize each of your test runs!

6.1.4 Running Test in Parallel

Test suites can take a long time to run. It's not uncommon for a full suite of unit tests to take tens of minutes to run in large software projects. By default, pytest runs all tests serially, in an undefined order. Since most computers have several CPUs, you can usually speed things up if you split the list of tests and run them on multiple CPUs.

To handle this approach, `pytest` provides the plugin `pytest-xdist`, which you can install with `pip`. Once installed, this plugin extends the `pytest` command line with the `--numprocesses` argument (shortened as `-n`), which accepts as its argument the number of CPUs to use. Running `pytest -n 4` would run your test suite using four parallel processes, balancing the load across the available CPUs.

Because the number of CPUs can change from one computer to another, the plugin also accepts the `auto` keyword as a value: in this case, it will probe the machine to retrieve the number of CPUs available and start this number of processes.

6.1.5 Creating Objects Used in Tests with fixtures

In unit testing, you'll often need to execute a set of common instructions before and after running a test, and those instructions will use certain components. For example, you might need an object that represents the configuration state of your application, and you'll likely want that object to be initialized before each test, then reset to its default values when the test is achieved. Same goes if your test relies on the temporary creation of a file: the file must be created before the test starts and deleted once the test is done. These components, known as *fixtures*, are set up before a test and cleaned up after the test has finished.

With `pytest`, fixtures are defined as simple functions. The fixture function should return the desired object(s) so that a test using that fixture can use that object.

Here's a pretty simple fixture:

```
1 import pytest
2
3 @pytest.fixture
4 def database():
5     return <some database connection>
6
7 def test_insert(database):
```

```
8 database.insert(123)
```

The database fixture is automatically used by any test that has database in its argument list. The `test_insert()` function will therefore receive as its first argument the result of the `database()` function and use that result as it wants. By using a fixture this way, we conveniently don't need to repeat the database initialization code several times.

Another common feature of code testing is tearing down after a test has used a fixture; for example, you may need to close a database connection. Implementing the fixture as a generator allows us to add teardown functionality, as shown in Example 6.5.

Example 6.5 Teardown functionality

```
1 import pytest
2
3 @pytest.fixture
4 def database():
5     db = <some database connection>
6     yield db
7     db.close()
8
9 def test_insert(database):
10     database.insert(123)
```

Because we used the `yield` keyword and made `database` a generator, the code after the `yield` statement will be run when the test is done. That code will close the database connection at the end of the test.

However, closing a database connection for each test might impose an unnecessary runtime cost, as tests may be able to reuse that same connection. In that case, you can pass the `scope` argument to the fixture decorator, specifying the scope of the fixture:

```
1 import pytest
```

6.1. THE BASICS OF TESTING

```
2
3 @pytest.fixture(scope="module")
4 def database():
5     db = <some database connection> yield db
6     db.close()
7
8 def test_insert(database):
9     database.insert(123)
```

By specifying the `scope="module"` parameter, you initialize the fixture once for the whole module, and the same database connection will be passed to all test functions requesting a database connection.

Finally, you can run some common code before and after your tests by marking fixtures as *automatically* used with the `autouse` keyword, rather than specifying them as an argument for each of the test functions. Specifying the `autouse=True` keyword argument to the `pytest.fixture()` function will make sure the fixture is called before running any test in the module or class it is defined in, as in this example:

```
1 import os
2
3 import pytest
4
5 @pytest.fixture(autouse=True)
6 def change_user_env():
7     curuser = os.environ.get("USER")
8     os.environ["USER"] = "foobar"
9     yield
10    os.environ["USER"] = curuser
11
12 def test_user():
13    assert os.getenv("USER") == "foobar"
```

Such automatically enabled features are handy, but make sure not to abuse fixtures: they are run before each and every test covered by their scope, so they can slow down a test run significantly.

6.1.6 Running Test Scenarios

When unit testing, it's common to need to run a set of tests against different versions of an object. You may want to run the same error-handling test with a bunch of different objects that trigger that error, or you may want to run an entire test suite against different drivers.

This latter case is one that we relied on heavily when developing *Gnocchi*, a time series database. Gnocchi provides an abstract class that we call the storage API. Any Python class can implement this abstract base and register itself to become a driver. The software loads the configured storage driver when required and uses the implemented storage API to store or retrieve data. In this case, we need a class of unit tests that runs against each driver—thus running against each implementation of this storage API — to be sure all drivers conform to what the callers expect.

An easy way to achieve this is by using *parameterized* fixtures, which will run all the tests that use them several times, once for each of the defined parameters. Example 6.6 shows an example of using parameterized fixtures to run a single test twice with different parameters: once for `mysql` and once for `postgresql`.

Example 6.6 Running a test using parameterized fixtures

```

1 import pytest
2
3 import myapp
4
5 @pytest.fixture(params=["mysql", "postgresql"])
6 def database(request):
7     d = myapp.driver(request.param)
8     d.start()
9     yield d
10    d.stop()
11
12 def test_insert(database):
13     database.insert("somedata")

```

In Example 6.6, the driver fixture is parameterized with two different values, each the name of a database driver that is supported by the application. When `test_insert` is run, it is actually run twice: once with a MySQL database connection and once with a PostgreSQL database connection. This allows us to easily reuse the same test with different scenarios, without adding many lines of code.

6.1.7 Controlled Test Using Mocking

Mock objects are simulated objects that mimic the behavior of real application objects, but in particular and controlled ways. These are especially useful in creating environments that describe precisely the conditions for which you would like to test code. You can replace all objects but the one with mock objects to isolate the behavior of your focus object.

One use case is in writing an HTTP client, for which it is likely impossible (or at least extremely complicated) to spawn the HTTP server and test it through all scenarios to return every possible value. HTTP clients are especially difficult to test for all failure scenarios.

A much simpler option is to build a set of mock objects that model these scenarios and to use them as an environment for testing your code.

The standard library for creating mock objects in Python is `mock`. Starting with Python 3.3, `mock` has been merged into the Python standard library as `unittest.mock`. You can, therefore, use a snippet like the following to maintain backward compatibility between Python 3.3 and earlier versions:

```
1 try:
2     from unittest import mock
3 except ImportError:
4     import mock
```

The `mock` library is pretty simple to use. Any attribute accessed on a mock. Mock

object is dynamically created at runtime. Any value can be set to such an attribute. Example 6.7 shows `mock` being used to create a fake object with a fake attribute.

Example 6.7 Accessing the `mock.Mock` attribute

```
1 >>> from unittest import mock
2 >>> m = mock.Mock()
3 >>> m.some_attribute = "hello world" >>> m.some_attribute
4 "hello world"
```

The ability to dynamically create attributes makes it easy to create the objects needed to test particular cases.

You can also dynamically create a method on such a malleable object, as in Example 6.8 where we create a fake method that always returns 42 and accepts anything as an argument.

Example 6.8 Creating methods on a `mock.Mock` object

```
1 >>> from unittest import mock
2 >>> m = mock.Mock()
3 >>> m.some_method.return_value = 42 >>> m.some_method()
4 42
5 >>> m.some_method("with", "arguments")
6 42
```

In just a few lines, your `mock.Mock` object now has a `some_method()` method that returns 42. It accepts any kind of argument, and there is no check on what the values are — yet.

Dynamically created methods can also have (intentional) side effects. Rather than being boilerplate methods that just return a value, they can be defined to execute useful code.

Example 6.9 creates a fake method that has the side effect of printing the “hello world” string.

Example 6.9 Creating methods on a mock.Mock object with side effects

```

1 >>> from unittest import mock >>> m = mock.Mock()
2 >>> def print_hello():
3 ...     print("hello world!")
4 ...     return 43
5 ...
6 >>> m.some_method.side_effect = print_hello
7 >>> m.some_method() hello world!
8 43
9 >>> m.some_method.call_count
10 1

```

We assign an entire function to the `some_method` attribute. This technique allows us to implement more complex scenarios in a test because we can plug any code needed for testing into a mock object. We then just need to pass this mock object to whichever function expects it.

The `call_count` attribute is a simple way of checking the number of times a method has been called.

The mock library uses the action/assertion pattern: this means that once your test has run, it's up to you to check that the actions you are mocking were correctly executed. Example 6.10 applies the `assert()` method to our mock objects to perform these checks.

Example 6.10 Checking method calls

```

1 >>> from unittest import mock
2 >>> m = mock.Mock()
3 >>> m.some_method('foo', 'bar')
4 <Mock name='mock.some_method()' id='26144272'>
5 >>> m.some_method.assert_called_once_with('foo', 'bar')
6 >>> m.some_method.assert_called_once_with('foo', mock.ANY)
7 >>> m.some_method.assert_called_once_with('foo', 'baz')
8 Traceback (most recent call last):

```

6.1. THE BASICS OF TESTING

```
9 File "<stdin>", line 1, in <module>
10 File "/usr/lib/python2.7/dist-packages/mock.py", line 846, in ←
    assert_called_once_with
11     return self.assert_called_with(*args, **kwargs)
12 File "/usr/lib/python2.7/dist-packages/mock.py", line 835, in ←
    assert_called_with
13     raise AssertionError(msg)
14 AssertionError: Expected call: some_method('foo', 'baz')
15 Actual call: some_method('foo', 'bar')
```

We create a method with the arguments `foo` and `bar` to stand in as our tests by calling the method. The usual way to check calls to a mock object is to use the `assert_called()` methods, such as `assert_called_once_with()`. To these methods, you need to pass the values that you expect callers to use when calling your mock method. If the values passed are not the ones being used, then `mock` raises an `AssertionError`. If you don't know what arguments may be passed, you can use `mock.ANY` as a value; that will match any argument passed to your mock method.

Another function of `mock` is in patching some function, method, or object from an external module, for which it provides a set of patching functions. In Example 6.11, we replace the `os.unlink()` function with a fake function we provide.

Example 6.11 Using `mock.patch`

```
1 >>> from unittest import mock
2 >>> import os
3 >>> def fake_os_unlink(path):
4 ...     raise IOError("Testing!") ...
5 >>> with mock.patch('os.unlink', fake_os_unlink)
6 ...     os.unlink('foobar')
7 ...
8 Traceback (most recent call last):
9   File "<stdin>", line 2, in <module>
10   File "<stdin>", line 2, in fake_os_unlink
11 IOError: Testing!
```

6.1. THE BASICS OF TESTING

When used as a context manager, `mock.patch()` replaces the target function with the function we provide so the code executed inside the context uses that patched method. With the `mock.patch()` method, it's possible to change any part of an external piece of code, making it behave in the required way in order to test all conditions in your application, as shown in Example 6.12.

Example 6.12 Using `mock.patch` to test a set of behaviour

```
1 from unittest import mock
2
3 import pytest
4 import requests
5
6
7 class WhereIsPythonError(Exception):
8     pass
9
10
11 def is_python_still_a_programming_language():
12     try:
13         r = requests.get("http://python.org")
14     except IOError:
15         pass
16     else:
17         if r.status_code == 200:
18             return 'Python is a programming language' in r.content
19         raise WhereIsPythonError("Something bad happened")
20
21
22 def get_fake_get(status_code, content):
23     m = mock.Mock()
24     m.status_code = status_code
25     m.content = content
26
27     def fake_get(url):
28         return m
29
```

```

30     return fake_get
31
32
33 def raise_get(url):
34     raise IOError("Unable to fetch url %s" % url)
35
36
37 @mock.patch('requests.get', get_fake_get(
38     200, 'Python is a programming language for sure'))
39 def test_python_is():
40     assert is_python_still_a_programming_language() is True
41
42
43 @mock.patch('requests.get', get_fake_get(
44     200, 'Python is no more a programming language'))
45 def test_python_is_not():
46     assert is_python_still_a_programming_language() is False
47
48
49 @mock.patch('requests.get', get_fake_get(404, 'Whatever'))
50 def test_bad_status_code():
51     with pytest.raises(WhereIsPythonError):
52         is_python_still_a_programming_language()
53
54
55 @mock.patch('requests.get', raise_get)
56 def test_ioerror():
57     with pytest.raises(WhereIsPythonError):
58         is_python_still_a_programming_language()

```

Example 6.12 implements a test suite that searches for all instances of the string “Python is a programming language” on the <http://python.org/> web page. There is no way to test negative scenarios (where this sentence is not on the web page) without modifying the page itself — something we’re not able to do, obviously. In this case, we’re using mock to cheat and change the behavior of the request so it returns a mocked reply with a fake page that doesn’t contain that string. This allows us to test

the negative scenario in which <http://python.org/> does not contain this sentence, making sure the program handles that case correctly.

This example uses the decorator version of `mock.patch()`. Using the decorator does not change the mocking behavior, but it is simpler when you need to use mocking within the context of an entire test function.

Using mocking, we can simulate any problem, such as a web server returning a 404 error, an I/O error, or a network latency issue. We can make sure code returns the correct values or raises the correct exception in every case, ensuring our code always behaves as expected.

6.1.8 Revealing Untested Code with coverage

A great complement to unit testing, the coverage tool identifies whether any of your code has been missed during testing. It uses code analysis tools and tracing hooks to determine which lines of your code have been executed; when used during a unit test run, it can show you which parts of your code base have been crossed over and which parts have not. Writing tests is useful, but having a way to know what part of your code you may have missed during the testing process is the cherry on the cake.

Install the coverage Python module on your system via `pip` to have access to the coverage program command from your shell.



Note

The command may also be named `python-coverage`, if you install coverage through your operating system installation software. This is the case on Debian, for example.

Using coverage in stand-alone mode is straightforward. It can be useful for showing you parts of your programs that are never run and which code might be

“dead code,” that is, code that could be removed without modifying the normal workflow of the program. The ability for coverage to pinpoint code that is never executed allows it to show which parts of the code are not being tested. All the test tools we’ve talked about so far in this chapter are integrated with coverage.

When using pytest, just install the `pytest-cov` plugin via `pip install pytest-pycov` and add a few option switches to generate a detailed code coverage output, as shown in Example 6.13.

Example 6.13 Using coverage with pytest

```
1  $ pytest --cov=gnocchiclient gnocchiclient/tests/unit
2  ----- coverage: platform darwin, python 3.6.4-final-0 ↵
   -----
3  Name                               Stmts Cover
4  -----
5  gnocchiclient/__init__.py          0 100%
6  gnocchiclient/auth.py              51 49%
7  gnocchiclient/benchmark.py        175 0%
8  [...]
9  -----
10 TOTAL                             2040 8%
11 === passed in 5.00 seconds ===
```

The `--cov` option enables the coverage report at the end of the test run. You need to pass the package name as an argument for the plugin to filter the coverage report properly. The output includes the lines of code that were not run and therefore have no tests. All you need to do now is spawn your favorite text editor and start writing tests for that code.

However, coverage goes one better, allowing you to generate nice HTML reports. Simply add the `--cov-report=html` flag, and the `htmlcov` directory from which you ran the command will be populated with HTML pages. Each page will show you which parts of your source code were or were not run.

If you want to be that guy, you can use the option `--cover-fail-under=COVE`

6.1. THE BASICS OF TESTING

R_MIN_PERCENTAGE, which will make the test suite fail if a minimum percentage of the code is not executed when the test suite is run. While having a good coverage percentage is a decent goal, and while the tool is useful to gain insight into the state of your test coverage, defining an arbitrary percentage value does not provide much in the way of insight.

6.1. THE BASICS OF TESTING

Coverage for **ceilometer.publisher** : 75%

12 statements

```
1  # -*- encoding: utf-8 -*-
2  #
3  # Copyright © 2013 Intel Corp.
4  # Copyright © 2013 eNovance
5  #
6  # Author: Yunhong Jiang <yunhong.jiang@intel.com>
7  #       Julien Danjou <julien@danjou.info>
8  #
9  # Licensed under the Apache License, Version 2.0 (the "License"); you may
10 # not use this file except in compliance with the License. You may obtain
11 # a copy of the License at
12 #
13 #     http://www.apache.org/licenses/LICENSE-2.0
14 #
15 # Unless required by applicable law or agreed to in writing, software
16 # distributed under the License is distributed on an "AS IS" BASIS, WITHOUT
17 # WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the
18 # License for the specific language governing permissions and limitations
19 # under the License.
20
21 import abc
22 from stevedore import driver
23 from ceilometer.openstack.common import network_utils
24
25 def get_publisher(url, namespace='ceilometer.publisher'):
26     """Get publisher driver and load it.
27
28     :param URL: URL for the publisher
29     :param namespace: Namespace to use to look for drivers.
30     """
31     parse_result = network_utils.urlsplit(url)
32     loaded_driver = driver.DriverManager(namespace, parse_result.scheme)
33     return loaded_driver.driver(parse_result)
34
35 class PublisherBase(object):
36     """Base class for plugins that publish the sampler."""
37
38     __metaclass__ = abc.ABCMeta
39
40     def __init__(self, parsed_url):
41         pass
42
43     @abc.abstractmethod
44     def publish_samples(self, context, samples):
45         """Publish samples into final conduit."""
```

Figure 6.1: Coverage of `ceilometer.publisher`

For example, a code coverage score of 100 percent does not necessarily mean

the code is entirely tested and you can rest. It only proves that your whole code path has been run; there is no indication that every possible condition has been tested. So while 100 percent is a respectable goal, it does not indicate anything conclusive.

You should use coverage information to consolidate your test suite and add tests for any code that is currently not being run. This is important, as it facilitates later project maintenance and increases your code's overall quality.

6.2 Virtual Environments

Earlier we mentioned the danger that your tests may not capture the absence of dependencies. Any application of significant size inevitably depends on external libraries to provide features the application needs, but there are many ways external libraries might cause issues on your operating system. Here are a few:

- Your system does not have the library you need packaged.
- Your system does not have the right version of the library you need packaged.
- You need two different versions of the same library for two different applications.

Those problems can happen when you deploy your application or later on, while it's running. Upgrading a Python library installed via your system manager might break your application in a snap without warning, for reasons as simple as an API change in the library being used by the application.

The solution is for each application to use a library directory that contains all the application's dependencies. This directory is then used to load the needed Python modules rather than the system-installed ones.

Such a directory is known as a *virtual environment*.

6.2.1 Setting Up a Virtual Environment

The tool `virtualenv` handles virtual environments automatically for you. Until Python 3.2, you'll find it in the `virtualenv` package that you can install using `pip install virtualenv`. If you use Python 3.3 or later, it's available directly via Python under the ``venv`` name.

To use the module, load the module as the main program with a destination directory as its argument, like so:

```
1 $ python3 -m venv myvenv
2 $ ls foobar
3 bin      include  lib      pyvenv.cfg
```

Once run, `venv` creates a `lib/pythonX.Y` directory and uses it to install `pip` into the virtual environment, which will be useful to install further Python packages.

You can then activate the virtual environment by “sourcing” the `activate` command with the following on Posix systems:

```
1 $ source myvenv/bin/activate
```

Or on Windows systems:

```
1 > \myvenv\Scripts\activate
```

Once you do that, your shell prompt should appear prefixed by the name of your virtual environment. Executing `python` will call the version of Python that has been copied into the virtual environment. You can check that it's working by reading the `sys.path` variable and checking that it has your virtual environment directory as its first component.

You can stop and leave the virtual environment at any time by calling the `deactivate` command:

```
1 $ deactivate
```

6.2. VIRTUAL ENVIRONMENTS

That's it. Also note that you are not forced to run `activate` if you want to use the Python installed in your virtual environment just once. Calling the `python` binary will also work:

```
1 $ myenv/bin/python
```

Now, while we're in our activated virtual environment, we do not have access to any of the modules installed and available on the main system. That is kind of the point of using a virtual environment, but it does mean we probably need to install the packages we need. To do that, use the standard `pip` command to install each package, and the packages will install in the right place, without changing anything about your system:

```
1 $ source myenv/bin/activate
2 (myenv) $ pip install six
3 Downloading/unpacking six
4   Downloading six-1.4.1.tar.gz
5   Running setup.py egg_info for package six
6
7 Installing collected packages: six
8   Running setup.py install for six
9
10 Successfully installed six
11 Cleaning up...
```

Voilà! We can install all the libraries we need and then run our application from this virtual environment, without breaking our system. It's then easy to see how we can script this to automatize the installation of a virtual environment based on a list of dependencies, as in Example 6.14.

Example 6.14 Automatic virtual environment creation

```
1 virtualenv myappvenv
2 source myappvenv/bin/activate
3 pip install -r requirements.txt
4 deactivate
```

It can still be useful to have access to your system-installed packages, so `virtualenv` allows you to enable them when creating your virtual environment by passing the `--system-site-packages` flag to the `virtualenv` command.

Inside `myvenv`, you will find a `pyvenv.cfg`, the configuration file for this environment. It doesn't have a lot of configuration options by default. You should recognize `include-system-site-package`, whose purpose is the same as the `--system-site-packages` of `virtualenv` that we described earlier.

As you might guess, virtual environments are incredibly useful for automated runs of unit test suites. Their use is so widespread that a particular tool has been built to address it.

6.3 Using `virtualenv` with `tox`

One of the central uses of virtual environments is to provide a clean environment for running unit tests. It would be detrimental if you were under the impression that your tests were working, when in fact they were not, for example, respecting the dependency list.

One way to ensure you're accounting for all the dependencies would be to write a script to deploy a virtual environment, install `setuptools`, and then install all of the dependencies required for both your application/library runtime and unit tests. Luckily, this is such a popular use case that an application dedicated to this task has already been built: *tox*.

The `tox` management tool aims to automate and standardize how tests are run in Python. To that end, it provides everything needed to run an entire test suite in a clean virtual environment, while also installing your application to check that the installation works fine.

Before using `tox`, you need to provide a configuration file named `tox.ini` that should be placed in the root directory of your project, beside your `setup.py` file.

```
1 $ touch tox.ini
```

You can then run tox successfully:

```
1 % tox
2 GLOB sdist-make: /home/jd/project/setup.py
3 python create: /home/jd/project/.tox/python
4 python inst: /home/jd/project/.tox/dist/project-1.zip
5 ----- summary -----
6 python: commands succeeded
7 congratulations :)
```

In this instance, tox creates a virtual environment in `.tox/python` using the default Python version. It uses `setup.py` to create a distribution of your package, which it then installs inside this virtual environment. No commands are run, because we did not specify any in the configuration file. This alone is not particularly useful.

We can change this default behavior by adding a command to run inside our test environment. Edit `tox.ini` to include the following:

```
1 [testenv]
2 commands=pytest
```

Now tox runs the command `pytest`. However, since we do not have `pytest` installed in the virtual environment, this command will likely fail. We need to list `pytest` as a dependency to be installed:

```
1 [testenv]
2 deps=pytest
3 commands=pytest
```

When run now, tox re-creates the environment, installs the new dependency, and runs the command `pytest`, which executes all of the unit tests. To add more dependencies, you can either list them in the `deps` configuration option, as is done here, or use the `-rfile` syntax to read from a file.

6.3.1 Re-creating an Environment

Sometimes you'll need to re-create an environment to, for example, ensure things work as expected when a new developer clones the source code repository and runs `tox` for the first time. For this, `tox` accepts a `--recreate` option that will rebuild the virtual environment from scratch based on parameters you lay out.

You define the parameters for all virtual environments managed by `tox` in the `[testenv]` section of `tox.ini`. And, as mentioned previously, `tox` can manage multiple Python virtual environments—indeed, it is possible to run our tests under a Python version other than the default one by passing the `-e` flag to `tox`, like so:

```

1  % tox -e py26
2  GLOB sdist-make: /home/jd/project/setup.py
3  py26 create: /home/jd/project/.tox/py26
4  py26 installdeps: nose
5  py26 inst: /home/jd/project/.tox/dist/rebuildd-1.zip
6  py26 runtests: commands[0] | nosetests
7  .....
8
9  -----
10 Ran 7 tests in 0.029s
11
12 OK
13
14 ----- summary -----
15 py26: commands succeeded
   congratulations :)
```

By default, `tox` simulate any environments that matches an existing Python version: `py24`, `py25`, `py26`, `py27`, `py30`, `py31`, `py32`, `py33`, `py34`, `py35`, `py36`, `jython` and `ppypy`! Furthermore, you can define your own environments. You just need to add another section named `[testenv:_envname_]`. If you want to run a particular command for just one of the environments, you can do so easily by listing the following in the `tox.ini` file:

```

1 [testenv]
2 deps=pytest
3 commands=pytest
4
5 [testenv:py36-coverage]
6 deps={[testenv]deps}
7     pytest-cov
8 commands=pytest --cov=myproject

```

By using `pytest --cov=myproject` under the `py36-coverage` section as shown here, you override the commands for the `py36-coverage` environment, meaning when you run `tox -e py36-coverage`, `pytest` is installed as part of the dependencies, but the command `pytest` is actually run instead with the coverage option. For that to work, the `pytest-cov` extension must be installed: to this end, we replace the `deps` value with the `deps` from `testenv` and add the `pytest-cov` dependency. Variable interpolation is also supported by `tox`, so you can refer to any other field from the `tox.ini` file and use it as a variable, the syntax being `{[env_name]variable_name}`. This allows us to avoid repeating the same things over and over again.

6.3.2 Using Different Python Versions

We can also create a new environment with an unsupported version of Python right away with the following in `tox.ini`:

```

1 [testenv]
2 deps=pytest
3 commands=pytest
4
5 [testenv:py21]
6 basepython=python2.1

```

When we run this, it will now (attempt to) use Python 2.1 to run the test suite — although since it is very unlikely you have this very ancient Python version installed on your system, I doubt this would work for you!

It's likely that you'll want to support multiple Python versions, in which case it would be useful to have `tox` run all the tests for all the Python versions you want to support by default. You can do this by specifying the environment list you want to use when `tox` is run without arguments:

```

1 [tox]
2 envlist=py35,py36,pypy
3
4 [testenv]
5 deps=pytest
6 commands=pytest

```

When `tox` is launched without any further arguments, all four environments listed are created, populated with the dependencies and the application, and then run with the command `pytest`.

6.3.3 Integrating Other Tests

We can also use `tox` to integrate tests like `flake8`, as discussed in Section 1.4. The following `tox.ini` file provides a `pep8` environment that will install `flake8` and run it:

```

1 [tox]
2 envlist=py35,py36,pypy,pep8
3
4 [testenv]
5 deps=pytest
6 commands=pytest
7
8 [testenv:pep8]
9 deps=flake8
10 commands=flake8

```

In this case, the `pep8` environment is run using the default version of Python, which is probably fine, though you can still specify the `basepython` option if you want to change that.

Tip

When running `tox`, you'll notice that all the environments are built and run sequentially. This can make the process very long, but since virtual environments are isolated, nothing might prevent you from running `tox` commands in parallel. This is exactly what the `detox` package does, by providing a `detox` command that runs all of the default environments from `envlist` in parallel. You should `pip install` it!

6.4 Testing Policy

Embedding testing code in your project is an excellent idea, but how that code is run is also extremely important. Too many projects have test code lying around that fails to run for some reason or other. This topic is not strictly limited to Python, but I consider it important enough to emphasize here: you should have a zero-tolerance policy regarding untested code. No code should be merged without a proper set of unit tests to cover it.

The minimum you should aim for is that each of the commits you push passes all the tests. Automating this process is even better. For example, OpenStack relies on a specific workflow based on Gerrit (a web-based code review service) and Zuul (a continuous integration and delivery service). Each commit pushed goes through the code review system provided by Gerrit, and Zuul is in charge of running a set of testing jobs. Zuul runs the unit tests and various higher-level functional tests for each project. This code review, which is executed by a couple of developers, makes sure all code committed has associated unit tests.

If you're using the popular `GitHub` hosting service, *Travis CI* is a tool that allows you to run tests after each push or merge or against pull requests that are submitted. While it is unfortunate that this testing is done post-push, it's still a fantastic way to

track regressions. Travis supports all significant Python versions out of the box, and it can be customized significantly. Once you've activated Travis on your project via the web interface, just add a `.travis.yml` file that will determine how the tests are run. Example 6.15 shows an example of a `travis.yml` file.

Example 6.15 A `.travis.yml` example file

```
1 language: python
2 python:
3   - "2.7"
4   - "3.6"
5 # command to install dependencies
6 install: "pip install -r requirements.txt --use-mirrors"
7 # command to run tests
8 script: pytest
```

With this file in place in your code repository and Travis enabled, the latter will spawn a set of jobs to test your code with the associated unit tests.

It's easy to see how you would customize this by simply adding dependencies and tests. Travis is a paid-for service, but the good news is that for open source projects, it's entirely free!

The **tox-travis** package is also worth looking into, as it will polish the integration between tox and Travis by running the correct tox target depending on the Travis environment being used. Example 6.16 shows an example of a `.travis.yml` file that will install `tox-travis` before running tox.

Example 6.16 A `.travis.yml` example file with *tox-travis*

```
1 sudo: false
2 language: python
3 python:
4   - "2.7"
5   - "3.6"
6 install: pip install tox-travis
```

```
7 script: tox
```

Using `tox-travis`, you can simply call `tox` as the script on Travis, and it will call `tox` with the environment you specify here in the `.travis.yml` file, building the necessary virtual environment, installing the dependency, and running the commands you specified in `tox.ini`. This makes it easy to use the same workflow both on your local development machine and on the Travis continuous integration platform.

These days, wherever your code is hosted, it is always possible to apply some automatic testing of your software and to make sure your project is moving forward, not being held back by the addition of bugs.



Tip

If your tests pass (e.g., Travis validates them) and the project owners approve the pull request, there's nothing preventing the pull request to be merged automatically. You can leverage [Mergify](#) to set that as an automatic process.

6.5 Robert Collins on Testing

Robert Collins is, among other things, the original author of the Bazaar distributed version control system. Today, he is a Distinguished Technologist at HP Cloud Services, where he works on OpenStack. Robert is also the author of many of the Python tools described in this book, such as *fixtures*, *testscenarios*, *testrepository*, and even *python-subunit* — you may have used one of his programs without knowing it!

“



What kind of testing policy would you advise using? Is it ever acceptable not to test code?

I think testing is an engineering trade-off: you must consider the likelihood of a failure slipping through to production undetected, the cost and size of an undetected failure, and cohesion of the team doing the work. Take OpenStack, which has 1,600 contributors: it's difficult to work with a nuanced policy with so many people with their own opinions. Generally speaking, a project needs some automated testing to check that the code will do what it is intended to do and that what it is intended to do is what is needed. Often that requires functional tests that might be in different code bases. Unit tests are excellent for speed and pinning down corner cases. I think it is okay to vary the balance between styles of testing, as long as there is testing.

Where the cost of testing is very high and the returns are very low, I think it's fine to make an informed decision not to test, but that situation is relatively rare: most things can be tested reasonably cheaply, and the benefit of catching errors early is usually quite high.

What are the best strategies when writing Python code to make testing manageable and improve the quality of the code?

Separate out concerns and don't do multiple things in one place; this makes reuse natural, and that makes it easier to put test doubles in place. Take a purely functional approach when possible; for example, in a single method either calculate something or change some state, but avoid doing both. That way you can test all of the calculating behaviors without dealing with state changes, such as writing to a database or talking to an HTTP server. The benefit works the other way around too—you can replace the calculation logic for tests to provoke corner case behavior and use mocks and test doubles to check

that the expected state propagation happens as desired. The most heinous things to test are deeply layered stacks with complex cross-layer behavioral dependencies. There you want to evolve the code so that the contract between layers is simple, predictable, and — most usefully for testing — replaceable.

What's the best way to organize unit tests in source code?

Have a clear hierarchy, like `$ROOT/$PACKAGE/tests`. I tend to do just one hierarchy for a whole source tree, for example `$ROOT/$PACKAGE/$SUBPACKAGE/tests`.

Within tests, I often mirror the structure of the rest of the source tree: `$ROOT/$PACKAGE/foo.py` would be tested in `$ROOT/$PACKAGE/tests/test_foo.py`.

The rest of the tree should not import from the tests tree, except perhaps in the case of a `test_suite/load_tests` function in the top level `__init__`. This permits you to easily detach the tests for small-footprint installations.

What do you see as the future of unit-testing libraries and frameworks in Python?

The significant challenges I see are these:

- The continued expansion of parallel capabilities in new machines, like phones with four CPUs. Existing unit test internal APIs are not optimized for parallel workloads. My work on the `StreamResult` Java class is aimed directly at resolving this.
- More complex scheduling support — a less ugly solution for the problems that class and module-scoped setup aim at.
- Finding some way to consolidate the vast variety of frameworks we have today: for integration testing, it would be great to be able to

6.5. ROBERT COLLINS ON TESTING

get a consolidated view across multiple projects that have different test runners in use.

”

CHAPTER 7



Methods and Decorators

Python's decorators are a handy way to modify functions. Decorators were first introduced in Python 2.2, with the `classmethod()` and `staticmethod()` decorators, but were overhauled to become more flexible and readable. Along with these two original decorators, Python now provides a few right out of the box and supports the simple creation of custom decorators. But it seems as though most developers do not understand how they work behind the scenes.

This chapter aims to change that — we'll cover what a decorator is and how to use it, as well as how to create your own decorators. Then we'll look at using decorators to create static, class, and abstract methods and take a close look at the `super` function, which allows you to place implementable code inside an abstract method.

7.1 Decorators and When to Use Them

A decorator is a function that takes another function as an argument and replaces it with a new, modified function. The primary use case for decorators is in factoring common code that needs to be called before, after, or around multiple functions. If

you’ve ever written Emacs Lisp code, you may have used the `defadvice` decorator, which allows you to define code called around a function. If you’ve used method combinations in the Common Lisp Object System (CLOS), Python decorators follow the same concepts. We’ll look at some simple decorator definitions, and then we’ll examine some common situations in which you’d use decorators.

7.1.1 Creating Decorators

The odds are good that you’ve already used decorators to make your own wrapper functions. The dullest possible decorator and simplest example is the `identity()` function, which does nothing except return the original function. Here is its definition:

```
1 def identity(f):  
2     return f
```

You would then use your decorator like this:

```
1 @identity  
2 def foo():  
3     return 'bar'
```

Here we give the decorator preceded by an `@symbol` and then the function we want to use it on. This is the same as writing:

```
1 def foo():  
2     return 'bar'  
3  
4 foo = identity(foo)
```

This decorator is useless and does nothing, but it works. Let’s look at another, more useful example in [Example 7.1](#).

Example 7.1 A decorator to organize functions in a dictionary

```
1 _functions = {}
2 def register(f):
3     global _functions
4     _functions[f.__name__] = f
5     return f
6
7 @register
8 def foo():
9     return 'bar'
```

In Example 7.1, the register decorator stores the decorated function name into a dictionary. The `_functions` dictionary can then be used and accessed using the function name to retrieve a function: `_functions['foo']` points to the `foo` function. In the following sections, I will explain how to write your own decorators. Then I'll cover how the built-in decorators provided by Python work and explain how (and when) to use them.

7.1.2 Writing Decorators

As mentioned, decorators are often used when refactoring repeated code around functions. Consider the following set of functions that need to check whether the username that they receive as an argument is the admin or not and, if the user is not an admin, raise an exception:

```
1 class Store(object):
2     def get_food(self, username, food):
3         if username != 'admin':
4             raise Exception("This user is not allowed to get food")
5         return self.storage.get(food)
6
7     def put_food(self, username, food):
8         if username != 'admin':
```

7.1. DECORATORS AND WHEN TO USE THEM

```
9         raise Exception("This user is not allowed to put food")
10    self.storage.put(food)
```

We can see there's some repeated code here. The obvious first step to making this code more efficient is to factor the code that checks for admin status:

```
1 def check_is_admin(username):
2     if username != 'admin':
3         raise Exception("This user is not allowed to get food")
4
5 class Store(object):
6     def get_food(self, username, food):
7         check_is_admin(username)
8         return self.storage.get(food)
9
10    def put_food(self, username, food):
11        check_is_admin(username)
12        self.storage.put(food)
```

We've moved the checking code into its own function. Now our code looks a bit cleaner, but we can do even better if we use a decorator, as shown in Example 7.2.

Example 7.2 Adding a decorator to the factored code

```
1 def check_is_admin(f):
2     def wrapper(*args, **kwargs):
3         if kwargs.get('username') != 'admin':
4             raise Exception("This user is not allowed to get food")
5         return f(*args, **kwargs)
6     return wrapper
7
8 class Store(object):
9     @check_is_admin
10    def get_food(self, username, food):
11        return self.storage.get(food)
12
13    @check_is_admin
```

```

14 def put_food(self, username, food):
15     self.storage.put(food)

```

We define our `check_is_admin` decorator and then call it whenever we need to check for access rights. The decorator inspects the arguments passed to the function using the `kwargs` variable and retrieves the `username` argument, performing the username check before calling the actual function. Using decorators like this makes it easier to manage common functionality. To anyone with much Python experience, this is probably old hat, but what you might not realize is that this naive approach to implementing decorators has some major drawbacks.

7.1.3 Stacking Decorators

You can also use several decorators on top of a single function or method, as shown in Example 7.3.

Example 7.3 Using more than one decorator with a single function

```

1 def check_user_is_not(username):
2     def user_check_decorator(f):
3         def wrapper(*args, **kwargs):
4             if kwargs.get('username') == username:
5                 raise Exception("This user is not allowed to get food")
6             return f(*args, **kwargs)
7         return wrapper
8     return user_check_decorator
9
10 class Store(object):
11     @check_user_is_not("admin")
12     @check_user_is_not("user123")
13     def get_food(self, username, food):
14         return self.storage.get(food)

```

7.1. DECORATORS AND WHEN TO USE THEM

Here, `check_user_is_not()` is a factory function for our decorator `user_check_decorator`. It creates a function decorator that depends on the `username` variable and then returns that variable. The function `user_check_decorator()` will serve as a function decorator for `get_food()`.

The function `get_food()` gets decorated twice using `check_user_is_not()`. The question here is in which username be checked first — `admin` or `user123`? The answer is in the following code, where I translated Example 7.3 into equivalent code without using a decorator.

```
1 class Store(object):
2     def get_food(self, username, food):
3         return self.storage.get(food)
4
5 Store.get_food = check_user_is_not("user123")(Store.get_food)
6 Store.get_food = check_user_is_not("admin")(Store.get_food)
```

The decorator list is applied from top to bottom, so the decorators closest to `def` keyword will be applied first and executed last. In the example above, the program will check for `admin` first and then for `user123`.

7.1.4 Writing Class Decorators

It's also possible to implement class decorators, though these are less often used in the wild. *Class decorators* work in the same way as function decorators, but they act on classes rather than functions. The following is an example of a class decorator that sets attributes for two classes:

```
1 import uuid
2
3 def set_class_name_and_id(klass):
4     klass.name = str(klass)
5     klass.random_id = uuid.uuid4()
6     return klass
```

7.1. DECORATORS AND WHEN TO USE THEM

```
7
8 @set_class_name_and_id
9 class SomeClass(object):
10     pass
```

When the class is loaded and defined, it will set the name and random_id attributes, like so:

```
1 >>> SomeClass.name
2 "<class '__main__.SomeClass'>"
3 >>> SomeClass.random_id UUID('d244dc42-f0ca-451c-9670-732dc32417cd')
```

As with function decorators, this can be handy for factorizing common code that manipulates classes.

Another possible use for class decorators is to wrap a function or class with classes. For example, class decorators are often used for wrapping a function that's storing a state. The following example wraps the print() function to check how many times it has been called in a session:

```
1 class CountCalls(object):
2     def __init__(self, f):
3         self.f = f
4         self.called = 0
5     def __call__(self, *args, **kwargs):
6         self.called += 1
7         return self.f(*args, **kwargs)
8
9 @CountCalls
10 def print_hello():
11     print("hello")
```

We can then use this to check how many times the function print_hello() has been called:

```
1 >>> print_hello.called
2 0
```



```

3 >>> print_hello()
4 hello
5 >>> print_hello.called
6 1

```

7.1.5 Retrieving Original Attributes with the `update_wrapper` Decorator

As mentioned before, a decorator replaces the original function with a new one built on the fly. However, this new function lacks many of the attributes of the original function, such as its docstring and its name. Example 7.4 shows how the function `foobar` loses its docstring and its name attribute once it is decorated with the `is_admin` decorator.

Example 7.4 A decorated function loses its docstring and name attributes.

```

1 >>> def is_admin(f):
2 ...     def wrapper(*args, **kwargs):
3 ...         if kwargs.get('username') != 'admin':
4 ...             raise Exception("This user is not allowed to get food")
5 ...         return f(*args, **kwargs)
6 ...     return wrapper
7 ...
8 >>> def foobar(username="someone"):
9 ...     """Do crazy stuff."""
10 ...     pass
11 ...
12 >>> foobar.func_doc
13 'Do crazy stuff.'
14 >>> foobar.__name__
15 'foobar'
16 >>> @is_admin
17 ... def foobar(username="someone"):
18 ...     """Do crazy stuff."""

```

7.1. DECORATORS AND WHEN TO USE THEM

```
19 ...     pass
20 ...
21 >>> foobar.__doc__
22 >>> foobar.__name__
23 'wrapper'
```

Not having the correct docstring and name attribute for a function can be problematic in various situations, such as when generating the source code documentation. Fortunately, the `functools` module in the Python standard library solves this problem with the `update_wrapper()` function, which copies the attributes from the original function that were lost to the wrapper itself. The source code of `update_wrapper()` is shown in Example 7.5.

Example 7.5 Source code of `functools.update_wrapper`

```
1 WRAPPER_ASSIGNMENTS = ('__module__', '__name__', '__qualname__', ' ←
   __doc__',
2                           '__annotations__')
3 WRAPPER_UPDATES = ('__dict__',)
4 def update_wrapper(wrapper,
5                    wrapped,
6                    assigned = WRAPPER_ASSIGNMENTS,
7                    updated = WRAPPER_UPDATES):
8     for attr in assigned:
9         try:
10             value = getattr(wrapped, attr)
11         except AttributeError:
12             pass
13         else:
14             setattr(wrapper, attr, value)
15     for attr in updated:
16         getattr(wrapper, attr).update(getattr(wrapped, attr, {}))
17     # Issue #17482: set __wrapped__ last so we don't inadvertently ←
18     # from the wrapped function when updating __dict__
19     wrapper.__wrapped__ = wrapped
```

```

20 # Return the wrapper so this can be used as a decorator via ↵
    partial()
21 return wrapper

```

In Example 7.5, the `update_wrapper()` source code highlights which attributes are worth saving when wrapping a function with a decorator. By default, the `__name__` attribute, `__doc__` attribute, and some other attributes are copied. You can also personalize which attributes of a function are copied to the decorated function. When we use `update_wrapper()` to rewrite Example 7.4, things are much nicer:

```

1 >>> def foobar(username="someone"):
2     ...     """Do crazy stuff."""
3     ...     pass
4     ...
5 >>> foobar = functools.update_wrapper(is_admin, foobar)
6 >>> foobar.__name__
7 'foobar'
8 >>> foobar.__doc__
9 'Do crazy stuff.'

```

Now the `foobar` function has the correct name and docstring even when decorated by `is_admin`.

7.1.6 wraps: A Decorator for Decorators

It can get tedious to use `update_wrapper()` manually when creating decorators, so `functools` provides a decorator for decorators called `wraps`. Example 7.6 shows the `wraps` decorator in use.

Example 7.6 Using `functools.wraps`

```

1 import functools
2
3 def check_is_admin(f):
4     @functools.wraps(f)

```

```

5  def wrapper(*args, **kwargs):
6      if kwargs.get('username') != 'admin':
7          raise Exception("This user is not allowed to get food")
8      return f(*args, **kwargs)
9  return wrapper
10
11 class Store(object):
12     @check_is_admin
13     def get_food(self, username, food):
14         return self.storage.get(food)

```

With `functools.wrap`, the decorator function `check_is_admin()` that returns the wrapper function takes care of copying the docstring, name function, and other information from the function `f` passed as argument. Thus, the decorated function (`get_food()`, in this case) still sees its unchanged signature.

7.1.7 Extracting Relevant Information with `inspect`

In our examples so far, we have assumed that the decorated function will always have a `username` passed to it as a keyword argument, but that might not be the case. It might instead have a bunch of information from which we need to extract the `username` to check. With this in mind, we'll build a smarter version of our decorator that can look at the decorated function's arguments and pull out what it needs.

For this, Python has the `inspect` module, which allows us to retrieve a function's signature and operate on it, as shown in Example 7.7.

Example 7.7 Retrieving function arguments using `inspect`

```

1  import functools
2  import inspect
3
4  def check_is_admin(f):
5      @functools.wraps(f)

```

```

6 def wrapper(*args, **kwargs):
7     func_args = inspect.getcallargs(f, *args, **kwargs)
8     if func_args.get('username') != 'admin':
9         raise Exception("This user is not allowed to get food")
10    return f(*args, **kwargs)
11 return wrapper
12
13 @check_is_admin
14 def get_food(username, type='chocolate'):
15     return type + " nom nom nom!"

```

The function that does the heavy lifting here is `inspect.getcallargs()`, which returns a dictionary containing the names and values of the arguments as key-value pairs. In our example, this function returns `{'username': 'admin', 'type': 'chocolate'}`. That means that our decorator does not have to check whether the `username` parameter is a positional or a keyword argument; all the decorator has to do is look for `username` in the dictionary.

Using `functools.wraps` and the `inspect` module, you should be able to write any custom decorator that you would ever need. However, do not abuse the `inspect` module: while being able to guess what the function will accept as an argument sounds handy, this capability can be fragile, breaking easily when function signatures change. Decorators are a terrific way to implement the *Don't Repeat Yourself* mantra so cherished by developers.

7.2 How Methods Work in Python

Methods are pretty simple to use and understand, and you've likely just used them correctly without delving in much deeper than you needed to. But to understand what certain decorators do, you need to know how methods work behind the scenes.

A method is a function that is stored as a class attribute. Let's have a look at what happens when we try to access such an attribute directly:

7.2. HOW METHODS WORK IN PYTHON

```
1 >>> class Pizza(object):
2 ...     def __init__(self, size):
3 ...         self.size = size
4 ...     def get_size(self):
5 ...         return self.size
6 ...
7 >>> Pizza.get_size
8 <function Pizza.get_size at 0x7fdbfd1a8b90>
```

We are told that `get_size()` is a function — but why is that? The reason is that at this stage, `get_size()` is not tied to any particular object. Therefore, it is treated as a normal function. Python will raise an error if we try to call it directly, like so:

```
1 >>> Pizza.get_size()
2 Traceback (most recent call last):
3   File "<stdin>", line 1, in <module>
4   TypeError: get_size() missing 1 required positional argument: 'self'
```

Python complains that we have not provided the necessary `self` argument. Indeed, as it is not bound to any object, the `self` argument cannot be set automatically. However, we are able to use the `get_size()` function not only by passing an arbitrary instance of the class to the method if we want to but also by passing *any* object, as long as it has the properties that the method expects to find. Here's an example:

```
1 >>> Pizza.get_size(Pizza(42))
2 42
```

This call works, just as promised. It is, however, not very convenient: we have to refer to the class every time we want to call one of its methods.

So Python goes the extra mile for us by binding a class's methods to its instances. In other words, we can access `get_size()` from any `Pizza` instance, and, better still, Python will automatically pass the object itself to the method's `self` parameter, like so:

Example 7.8 Calling bound `get_size`

```

1 >>> Pizza(42).get_size
2 <bound method Pizza.get_size of <__main__.Pizza object at 0x7f3138827910>>
3 >>> Pizza(42).get_size()
4 42

```

As expected, we do not have to provide any argument to `get_size`, since it's a bound method: its `self` argument is automatically set to our `Pizza` instance. Here is an even clearer example:

```

1 >>> m = Pizza(42).get_size
2 >>> m()
3 42

```

As long as you have a reference to the bound method, you do not even have to keep a reference to your `Pizza` object. Moreover, if you have a reference to a method but you want to find out which object it is bound to, you can just check the method's `__self__` property, like so:

```

1 >>> m = Pizza(42).get_size
2 >>> m.__self__
3 <__main__.Pizza object at 0x7f3138827910>
4 >>> m == m.__self__.get_size
5 True

```

Obviously, we still have a reference to our object, and we can find it back if we want.

7.3 Static methods

Static methods belong to a class, rather than an instance of a class, so they don't actually operate on or affect class instances. Instead, a static method operates on the param-

7.3. STATIC METHODS

eters it takes. Static methods are generally used to create utility functions, because they do not depend on the state of the class or its objects.

For example, in Example 7.9, we have the static `mix_ingredients()` method that belongs to the `Pizza` class but could actually be used to mix ingredients for any other food.

Example 7.9 Creating a static method as part of a class

```
1 class Pizza(object):
2     @staticmethod
3     def mix_ingredients(x, y):
4         return x + y
5
6     def cook(self):
7         return self.mix_ingredients(self.cheese, self.vegetables)
```

You could write `mix_ingredients` as a non-static method if you wanted to, but it would take a `self` argument that would never actually be used. Using the `@staticmethod` decorator gives us several things.

The first is speed: Python does not have to instantiate a bound method for each `Pizza` object we create. Bound methods are objects, too, and creating them has a CPU and memory cost — even if it's low. Using a static method lets us avoid that, like so:

```
1 >>> Pizza().cook is Pizza().cook
2 False
3 >>> Pizza().mix_ingredients == Pizza.mix_ingredients
4 True
5 >>> Pizza().mix_ingredients == Pizza().mix_ingredients
6 True
```

Second, static methods improve the readability of the code. When we see `@staticmethod`, we know that the method does not depend on the state of the object.

Third, static methods can be overridden in subclasses. If instead of a static

method, we used a `mix_ingredients()` function defined at the top level of our module, a class inheriting from `Pizza` wouldn't be able to change the way we mix ingredients for our pizza without overriding the `cook` method itself. With static methods, the subclasses can override the method for their own purposes.

Unfortunately, Python is not always able to detect for itself whether a method is static or not — I call that a defect of the language design. One possible approach is to add a check that detects such pattern and emits a warning using `flake8`. We will look into how to do this in Section 9.3.

7.4 Class Methods

Class methods are bound to a class rather than its instances. That means that those methods cannot access the state of the object but only the state and methods of the class. Example 7.10 shows how to write a class method.

Example 7.10 A class method is bound to its class

```

1  >>> class Pizza(object):
2  ...     radius = 42
3  ...     @classmethod
4  ...     def get_radius(cls):
5  ...         return cls.radius
6  ...
7  >>> Pizza.get_radius
8  <bound method type.get_radius of <class '__main__.Pizza'>>
9  >>> Pizza().get_radius
10 <bound method type.get_radius of <class '__main__.Pizza'>>
11 >>> Pizza.get_radius == Pizza().get_radius
12 True
13 >>> Pizza.get_radius()
14 42

```

7.5. ABSTRACT METHODS

As you can see, there are various ways to access the `get_radius()` class method, but however you choose to access it, the method is always bound to the class it is attached to. Also, its first argument must be the class itself. Remember: classes are objects too!

Class methods are principally useful for creating *factory methods*, which instantiate objects using a different signature than `__init__`.

```
1 class Pizza(object):
2     def __init__(self, ingredients):
3         self.ingredients = ingredients
4
5     @classmethod
6     def from_fridge(cls, fridge):
7         return cls(fridge.get_cheese() + fridge.get_vegetables())
```

If we used a `@staticmethod` here instead of a `@classmethod`, we would have to hardcode the `Pizza` class name in our method, making any class inheriting from `Pizza` unable to use our factory for its own purposes. In this case, however, we provide a `from_fridge()` factory method that we can pass a `Fridge` object to. If we call this method with something like `Pizza.from_fridge(myfridge)`, it returns a brand-new `Pizza` with ingredients taken from what's available in `myfridge`.

Any time you write a method that cares only about the class of the object and not about the object's state, it should be declared as a class method.

7.5 Abstract Methods

An *abstract method* is defined in an abstract base class that may not itself provide any implementation. When a class has an abstract method, it cannot be instantiated. As a consequence, an *abstract class* (defined as a class that has at least one abstract method) must be used as a parent class by another class. This subclass will be in charge of implementing the abstract method, making it possible to instantiate the parent class.

7.5. ABSTRACT METHODS

We can use abstract base classes to make clear the relationships between other, connected classes derived from the base class but have the abstract base class itself be impossible to instantiate. By using abstract base classes, you can ensure the classes derived from the base class implement particular methods from the base class, or an exception will be raised. The following example shows the simplest way to write an abstract method in Python:

```
1 class Pizza(object):
2     @staticmethod
3     def get_radius():
4         raise NotImplementedError
```

With this definition, any class inheriting from `Pizza` must implement and override the `get_radius()` method; otherwise, calling the method raises the exception shown here. This is handy for making sure that each subclass of `Pizza` implements its own way of computing and returning its radius.

This way of implementing abstract methods has a drawback: if you write a class that inherits from `Pizza` but forget to implement `get_radius()`, the error is raised only if you try to use that method at runtime.

Example 7.11 Calling an abstract method

```
1 >>> Pizza()
2 <__main__.Pizza object at 0x7fb747353d90>
3 >>> Pizza().get_radius()
4 Traceback (most recent call last):
5   File "<stdin>", line 1, in <module>
6   File "<stdin>", line 3, in get_radius
7   NotImplementedError
```

As `Pizza` is directly instantiable, there's no way to prevent this from happening. One way to make sure you get an early warning about forgetting to implement and override the method, or trying to instantiate an object with abstract methods, is to use Python's built-in `abc` (abstract base classes) module instead, like so:

Example 7.12 Implementing an abstract method using abc

```

1 import abc
2
3 class BasePizza(object, metaclass=abc.ABCMeta):
4
5     @abc.abstractmethod
6     def get_radius(self):
7         """Method that should do something."""

```

The `abc` module provides a set of decorators to use on top of methods that will be defined as abstracts and a metaclass to enable this. When you use `abc` and its special metaclass, as shown above, instantiating a `BasePizza` or a class inheriting from it that doesn't override `get_radius()` causes a `TypeError`:

```

1 >>> BasePizza()
2 Traceback (most recent call last):
3   File "<stdin>", line 1, in <module>
4   TypeError: Can't instantiate abstract class BasePizza with abstract ↔
   methods get_radius

```

We try to instantiate the abstract `BasePizza` class and are immediately told it can't be done!

While using abstract methods doesn't guarantee that the method is implemented by the user, this decorator helps you catch the error earlier. This is especially handy when you are providing interfaces that must be implemented by other developers; it's a good documentation hint.

7.6 Mixing Static, Class, and Abstract Methods

Each of these decorators is useful on its own, but the time may come when you'll have to use them together.

7.6. MIXING STATIC, CLASS, AND ABSTRACT METHODS

For example, you could define a factory method as a class method while forcing the implementation to be made in a subclass. In that case, you'd need to have a class method defined as both an abstract method and a class method. This section gives some tips that will help you with that.

First, an abstract method's prototype is not set in stone. When you implement the method, there is nothing stopping you from extending the argument list as you see fit. Example 7.13 is an example of code in which a subclass extends the signature of the abstract method of its parent.

Example 7.13 Using a subclass to extend the signature of the abstract method of its parent

```
1 import abc
2
3 class BasePizza(object, metaclass=abc.ABCMeta):
4
5     @abc.abstractmethod
6     def get_ingredients(self):
7         """Returns the ingredient list."""
8
9 class Calzone(BasePizza):
10     def get_ingredients(self, with_egg=False):
11         egg = Egg() if with_egg else None
12         return self.ingredients + [egg]
```

We define the Calzone subclass to inherit from the BasePizza class. We can define the Calzone subclass's methods any way we like, as long as they support the interface we define in BasePizza. This includes implementing the methods as either class or static methods. The following code defines an abstract `get_ingredients()` method in the base class and a static `get_ingredients()` method in the DietPizza subclass:

```
1 import abc
2
3 class BasePizza(object, metaclass=abc.ABCMeta):
```

7.6. MIXING STATIC, CLASS, AND ABSTRACT METHODS

```
4
5 @abc.abstractmethod
6 def get_ingredients(self):
7     """Returns the ingredient list."""
8
9 class DietPizza(BasePizza):
10     @staticmethod
11     def get_ingredients():
12         return None
```

Even though our static `get_ingredients()` method doesn't return a result based on the object's state, it supports our abstract `BasePizza` class's interface, so it's still valid.

It is also possible to use the `@staticmethod` and `@classmethod` decorators on top of `@abstractmethod` in order to indicate that a method is, for example, both static and abstract, as shown in Example 7.14.

Example 7.14 Using class method decorator with an abstract method

```
1 import abc
2
3 class BasePizza(object, metaclass=abc.ABCMeta):
4
5     ingredients = ['cheese']
6
7     @classmethod
8     @abc.abstractmethod
9     def get_ingredients(cls):
10         """Returns the ingredient list."""
11         return cls.ingredients
```

The abstract method `get_ingredients()` needs to be implemented by a subclass, but it's also a class method, meaning the first argument it will receive will be a class (not an object).

Note that by defining `get_ingredients()` as a class method in `BasePizza` like

this, you are not forcing any subclasses to define `get_ingredients()` as a class method — it could be a regular method. The same would apply if we had defined it as a static method: there's no way to force subclasses to implement abstract methods as a specific kind of method. As we have seen previously, you can change the signature of an abstract method when implementing it in a subclass in any way you like.

7.7 Putting Implementations in Abstract Methods

Hold the fort: in Example 7.14, we have an implementation in an abstract method. Can we do that? The answer is yes. Python does not have a problem with it! You can put code in your abstract methods and call it using `super()`, as demonstrated in Example 7.15.

Example 7.15 Using an implementation in an abstract method

```

1 import abc
2
3 class BasePizza(object, metaclass=abc.ABCMeta):
4
5     default_ingredients = ['cheese']
6
7     @classmethod
8     @abc.abstractmethod
9     def get_ingredients(cls):
10         """Returns the default ingredient list."""
11         return cls.default_ingredients
12
13 class DietPizza(BasePizza):
14     def get_ingredients(self):
15         return [Egg()] + super(DietPizza, self).get_ingredients()

```

In this example, every `Pizza` you make that inherits from `BasePizza` has to override the `get_ingredients()` method, but every `Pizza` also has access to the base class's default mechanism for getting the ingredients list. This mechanism is especially useful when providing an interface to implement while also providing base code that might be useful to all inheriting classes.

7.8 The Truth About `super`

Python has always allowed developers to use both single and multiple inheritances to extend their classes, but even today, many developers do not seem to understand how these mechanisms, and the `super()` method that is associated with them, work. To fully understand your code, you need to understand the trade-offs.

Multiple inheritances are used in many places, particularly in code involving a *mixin* pattern. A *mixin* is a class that inherits from two or more other classes, combining their features.



Note

Many of the pros and cons of single and multiple inheritances, composition, or even duck typing are out of scope for this book, so we won't cover everything here. If you are not familiar with these notions, I suggest you read about them to form your own opinions.

As you should know by now, classes are objects in Python. The construct used to create a class is a special statement that you should be well familiar with: `class classname(expression of inheritance)`.

The code in parentheses is a Python expression that returns the list of class objects to be used as the class's parents. Ordinarily, you would specify them directly, but you could also write something like this to specify the list of parent objects:


```

1 >>> def parent():
2 ...     return object
3 ...
4 >>> class A(parent()):
5 ...     pass
6 ...
7 >>> A.mro()
8 [<class '__main__.A'>, <type 'object'>]

```

This code works as expected: we declare class `A` with `object` as its parent class. The class method `mro()` returns the method resolution order used to resolve attributes — it defines how the next method to call is found via the tree of inheritance between classes. The current MRO system was first implemented in Python 2.3, and its internal workings are described in the Python 2.3 release notes. It defines how the tree of inheritance between classes is browsed to find the method to call.

We already saw that the canonical way to call a method in a parent class is to use the `super()` function, but what you probably don't know is that `super()` is actually a constructor and you instantiate a `super` object each time you call it. It takes either one or two arguments: the first argument is a class, and the second, optional argument is either a subclass or an instance of the first argument.

The object returned by the constructor functions as a proxy for the parent classes of the first argument. It has its own `__getattr__` method that iterates over the classes in the MRO list and returns the first matching attribute it finds. The `__getattr__` method is called when an attribute of the `super()` object is retrieved, as shown in Example 7.16.

Example 7.16 The `super` function is a constructor that instantiates a `super` object.

```

1 >>> class A(object):
2 ...     bar = 42
3 ...     def foo(self):
4 ...         pass

```

```

5 ...
6 >>> class B(object):
7 ...     bar = 0
8 ...
9 >>> class C(A, B):
10 ...     xyz = 'abc'
11 ...
12 >>> C.mro()
13 [<class '__main__.C'>, <class '__main__.A'>, <class '__main__.B'>, < ←
    type 'object'>]
14 >>> super(C, C()).bar
15 42
16 >>> super(C, C()).foo
17 <bound method C.foo of <__main__.C object at 0x7f0299255a90>>
18 >>> super(B).__self__
19 >>> super(B, B()).__self__
20 <__main__.B object at 0x1096717f0>

```

When requesting an attribute of the super object of an instance of C, the `__getattribute__` method of the `super()` object walks through the MRO list and returns the attribute from the first class it finds that has the super attribute.

In Example 7.16, we called `super` with two arguments, meaning we used a *bound* super object. If we call `super()` with only one argument, it returns an *unbound* super object instead:

```

1 >>> super(C)
2 <super: <class 'C'>, NULL>

```

Since no instance has been provided as the second argument, the super object cannot be bound to any instance. Therefore, you cannot use this unbound object to access class attributes. If you try, you'll get the following errors:

```

1 >>> super(C).foo
2 Traceback (most recent call last):
3   File "<stdin>", line 1, in <module>
4 AttributeError: 'super' object has no attribute 'foo'

```

```

5 >>> super(C).bar
6 Traceback (most recent call last):
7   File "<stdin>", line 1, in <module>
8 AttributeError: 'super' object has no attribute 'bar'
9 >>> super(C).xyz
10 Traceback (most recent call last):
11   File "<stdin>", line 1, in <module>
12 AttributeError: 'super' object has no attribute 'xyz'

```

At first glance, it might seem like this unbound kind of super object is useless, but actually the way the super class implements the descriptor proto-col `__get__` makes unbound super objects useful as class attributes:

```

1 >>> class D(C):
2 ...     sup = super(C)
3 ...
4 >>> D().sup
5 <super: <class 'C'>, <D object>>
6 >>> D().sup.foo
7 <bound method D.foo of <__main__.D object at 0x7f0299255bd0>>
8 >>> D().sup.bar
9 42

```

The unbound super object's `__get__` method is called using the instance `super(C).__get__(D())` and the attribute name `foo` as arguments, allowing it to find and resolve `foo`.

Note



Even if you've never heard of the descriptor protocol, it's likely you've used it through the `@property` decorator without knowing it. The *descriptor protocol* is the mechanism in Python that allows an object stored as an attribute to return something other than itself. This protocol is not covered in this book, but you can find out more about it in [the Python data model documentation](#).

7.9. SUMMARY

There are plenty of situations in which using `super()` can be tricky, such as when handling different method signatures along the inheritance chain. Unfortunately, there's no silver bullet for all occasions. The best precaution is to use tricks such as having all your methods accept their arguments using `*args`, `**kwargs`.

Since Python 3, `super()` has picked up a bit of magic: it can now be called from within a method without any arguments. When no arguments are passed to `super()`, it automatically searches the stack frame for arguments:

```
1 class B(A):  
2     def foo(self):  
3         super().foo()
```

The standard way of accessing parent attributes in subclasses is `super`, and you should always use it. It allows cooperative calls of parent methods without any surprises, such as parent methods not being called or being called twice when multiple inheritances are used.

7.9 Summary

Equipped with what you learned in this chapter, you should be unbeatable on everything that concerns methods definition in Python. Decorators are essential when it comes to code factorization, and proper use of the built-in decorators provided by Python can vastly improve the neatness of your Python code. Abstract classes are especially useful when providing API to other developers and services.

Class inheritance is not often fully understood, and having an overview of the internal machinery of the language is a good way to fully apprehend how this works. There should be no secrets left on this topic for you now!

CHAPTER 8



Functional Programming

Many Python developers are unaware of the extent to which you can use functional programming in Python, which is a shame: with few exceptions, functional programming allows you to write more concise and efficient code. Moreover, Python's support for functional programming is extensive.

This chapter will cover some of the functional programming aspects of Python, including creating and using generators. You'll learn about the most useful functional packages and functions available and how to use them in combination to get the most efficient code.

Note

If you want to get serious about functional programming, here's my advice: take a break from Python and learn a hugely functional programming language, such as Lisp. I know it might sound strange to talk about Lisp in a Python book, but playing with Lisp for several years taught me how to "think functional." You may not develop the thought processes necessary to make full use of functional programming if all your experience comes from imperative and object-oriented programming. Lisp isn't purely functional itself, but it has more focus on functional programming than you'll find in Python.

8.1 Creating Pure Functions

When you write code using a functional style, your functions are designed to have no side effects: instead, they take an input and produce an output without keeping state or modifying anything not reflected in the return value. Functions that follow this ideal are referred to as purely functional.

Let's start with an example of a regular, non-pure function that removes the last item in a list:

Example 8.1 A non-pure function

```
1 def remove_last_item(mylist):
2     """Removes the last item from a list."""
3     mylist.pop(-1) # This modifies mylist
```

Example 8.2 A pure function

```
1 def butlast(mylist):
2     """Like butlast in Lisp; returns the list without the last element ←
   ."""
```

```
return mylist[: -1] # This returns a copy of mylist
```

We define a `butlast()` function to work like `butlast` in Lisp, in that it returns the list without the last element without modifying the original list. Instead, it returns a copy of the list that has the modifications in place, allowing us to keep the original. The practical advantages of using functional programming include the following:

- **Modularity.** Writing with a functional style forces a certain degree of separation in solving your individual problems and makes sections of code easier to reuse in other contexts. Since the function does not depend on any external variable or state, call it from a different piece of code is straightforward.
- **Brevity.** Functional programming is often less verbose than other paradigms.
- **Concurrency.** Purely functional functions are thread-safe and can run concurrently. Some functional languages do this automatically, which can be a big help if you ever need to scale your application, though this is not quite the case yet in Python.
- **Testability.** Testing a functional program is incredibly easy: all you need is a set of inputs and an expected set of outputs. They are idempotent, meaning that calling the same function over and over with the same arguments will always return the same result.

8.2 Generators

A *generator* is an object that behaves like an iterator, in that it generates and returns a value on each call of its `next()` method until a `StopIteration` is raised. Generators, first introduced in [PEP 255](#), offer an easy way to create objects that implement the [iterator protocol](#). While writing generators in a functional style is not strictly necessary, doing so makes them easier to write and debug, and in practice they often are written functionally.

8.2. GENERATORS

To create a generator, just write a regular Python function that contains a `yield` statement. Python will detect the use of `yield` and tag the function as a generator. When execution reaches the `yield` statement, the function returns a value as with a `return` statement, but with one notable difference: the interpreter will save a stack reference, and this will be used to resume the function's execution when the `next()` function is called again.

When functions are executed, the chaining of their execution produces a *stack* — function calls are said to be stacked on each other. When a function returns, it's removed from the stack, and the value it returns is passed to the calling function. In the case of a generator, the function does not really return but *yields* instead. Python therefore saves the state of the function as a stack reference, resuming the execution of the generator at the point it saved when the next iteration of the generator is needed.

8.2.1 Creating a Generator

As mentioned, you create a generator by writing a normal function and including `yield` in the function's body. Example 8.3 creates a generator called `mygenerator()` that includes three `yields`, meaning it will iterate with the next three calls to `next()`.

Example 8.3 Creating a generator

```
1 >> def mygenerator():
2 ...     yield 1
3 ...     yield 2
4 ...     yield 'a'
5 ...
6 >>> mygenerator()
7 <generator object mygenerator at 0x10d77fa50>
8 >>> g = mygenerator()
9 >>> next(g)
10 1
11 >>> next(g)
12 2
```


8.2. GENERATORS

```
13 >>> next(g)
14 'a'
15 >>> next(g)
16 Traceback (most recent call last):
17   File "<stdin>", line 1, in <module>
18 StopIteration
```

When it runs out of yield statements, `StopIteration` is raised at the next call to `next()`.

In Python, generators keep a reference to the stack when a function yields something, and they resume this stack when a call to `next()` is executed again.

The naive approach when iterating over any data without using generators is to build the entire list first, which often consumes memory wastefully.

Say we want to find the first number between 1 and 10,000,000 that's equal to 50,000. Sounds easy, doesn't it? Let's make this a challenge. We'll run Python with a memory constraint of 128MB and try the naive approach of first building the entire list:

```
1 $ ulimit -v 131072
2 $ python
3 >>> a = list(range(10000000))
4 Traceback (most recent call last):
5   File "<stdin>", line 1, in <module>
6 MemoryError
```

Uh-oh. Turns out we can't build a list of 10 million items with only 128MB of memory!



Note

In Python 3, `range()` returns a generator when iterated. To get a generator in Python 2, you have to use `xrange()` instead. This function doesn't exist in Python 3 anymore, since it's redundant.

Let's try using a generator instead, with the same 128MB restriction:

```
1 $ ulimit -v 131072
2 $ python3
3 >>> for value in range(10000000):
4 ...     if value == 50000:
5 ...         print("Found it")
6 ...         break
7 ...
8 Found it
```

This time, our program executes without issue. When it is iterated over, the `range()` class returns a generator that dynamically generates our list of integers. Better still, since we are only interested in the 50,000th number, instead of building the full list, the generator only had to generate 50,000 numbers before it stopped.

By generating values on the fly, generators allow you to handle large data sets with minimal consumption of memory and processing cycles. Whenever you need to work with a huge number of values, generators can help you handle them efficiently.

8.2.2 Returning and Passing Values with `yield`

A `yield` statement also has a less commonly used feature: it can return a value in the same way as a function call. This allows us to pass a value to a generator by calling its `send()` method. As an example of using `send()`, we'll write a function called `shorten()` that takes a list of strings and returns a list consisting of those same strings, only truncated (Example 8.4).

Example 8.4 `yield` returning a value

```
1 def shorten(string_list):
2     length = len(string_list[0])
3     for s in string_list:
4         length = yield s[:length]
5
```

8.2. GENERATORS

```
6 mystringlist = ['loremipsum', 'dolorsit', 'ametfoobar']
7 shortstringlist = shorten(mystringlist)
8 result = []
9 try:
10     s = next(shortstringlist)
11     result.append(s)
12     while True:
13         number_of_vowels = len(
14             list(filter(lambda letter: letter in 'aeiou', s))
15         )
16         # Truncate the next string depending
17         # on the number of vowels in the previous one
18         s = shortstringlist.send(number_of_vowels)
19         result.append(s)
20 except StopIteration:
21     pass
```

In this example, we've written a function called `shorten()` that takes a list of strings and returns a list consisting of those same strings, only truncated. The length of each truncated string is equal to the number of vowels in the previous string: `loremipsum` has four vowels, so the second value returned by the generator will be the first four letters of `dolorsit`; `dolo` has only two vowels, so `ametfoobar` will be truncated to its first two letters `am`. The generator then stops and raises `StopIteration`. Our generator thus returns:

```
1 ['loremipsum', 'dolo', 'am']
```

Using `yield` and `send()` in this fashion allows Python generators to function like *coroutines* seen in Lua and other languages.

PEP 289 introduced generator expressions, making it possible to build one-line generators using a syntax similar to list comprehension:

```
1 >>> (x.upper() for x in ['hello', 'world'])
2 <generator object <genexpr> at 0x7ffab3832fa0>
3 >>> gen = (x.upper() for x in ['hello', 'world'])
4 >>> list(gen)
```

```
5 ['HELLO', 'WORLD']
```

In this example, `gen` is a generator, just as if we had used the `yield` statement. The `yield` in this case is implicit.

8.2.3 Inspecting Generators

You can check whether a function is considered a generator or not yourself by using `inspect.isgeneratorfunction`. In Example 8.5, we create a simple generator and inspect it.

Example 8.5 Checking whether a function is a generator

```
1 >>> import inspect
2 >>> def mygenerator():
3 ...     yield 1
4 ...
5 >>> inspect.isgeneratorfunction(mygenerator)
6 True
7 >>> inspect.isgeneratorfunction(sum)
8 False
```

You need to import the `inspect` package to use the `isgeneratorfunction()`. Then just pass it the name of the function to inspect. Reading the source code of `inspect.isgeneratorfunction()` gives us some insight into how Python marks functions as being generators (see Example 8.6).

Example 8.6 Source code of `inspect.isgeneratorfunction`

```
1 def isgeneratorfunction(object):
2     """Return true if the object is a user-defined generator function.
3
4     Generator function objects provides same attributes as functions.
5
```

8.2. GENERATORS

```
6 See help(isfunction) for attributes listing."""
7 return bool((isFunction(object) or ismethod(object)) and
8             object.func_code.co_flags & CO_GENERATOR)
```

The `isgeneratorfunction()` checks that the object is a function or a method and that its code has the special `CO_GENERATOR` flag set. This example shows how easy it is to understand how Python works under the hood.

The `inspect` package provides another useful function, `inspect.getgenerators_tate()`, which gives the current state of the generator. We'll use it on `mygenerator()` here at different points of execution:

```
1 >>> import inspect
2 >>> def mygenerator():
3 ...     yield 1
4 ...
5 >>> gen = mygenerator()
6 >>> gen
7 <generator object mygenerator at 0x7f94b44fec30>
8 >>> inspect.getgeneratorstate(gen)
9 'GEN_CREATED'
10 >>> next(gen)
11 1
12 >>> inspect.getgeneratorstate(gen)
13 'GEN_SUSPENDED'
14 >>> next(gen)
15 Traceback (most recent call last):
16   File "<stdin>", line 1, in <module>
17 StopIteration
18 >>> inspect.getgeneratorstate(gen)
19 'GEN_CLOSED'
```

This allows us to determine whether the generator is waiting to be run for the first time (`GEN_CREATED`), currently being executed by the interpreter (`GEN_RUNNING`), waiting to be resumed by a call to `next()` (`GEN_SUSPENDED`), or finished running (`GEN_CLOSED`). This might come in handy to debug your generators.

8.3 List comprehensions

List comprehension, or *listcomp* for short, allows you to define a list's contents inline with its declaration. To make a list into a listcomp, you must wrap it in square brackets as usual, but also include an expression that will generate the items in the list and a for loop to loop through them.

The following example creates a list without using list comprehension:

```
1 >>> x = []
2 >>> for i in (1, 2, 3):
3 ...     x.append(i)
4 ...
5 >>> x
6 [1, 2, 3]
```

And this next example uses list comprehension to make the same list with a single line:

```
1 >>> x = [i for i in (1, 2, 3)]
2 >>> x
3 [1, 2, 3]
```

Using a list comprehension presents two advantages: code written using listcomps is usually shorter and therefore compiles down to fewer operations for Python to perform. Rather than creating a list and calling `append` over and over, Python can just create the list of items and move them into a new list in a single operation.

You can use multiple for statements together and use if statements to filter out items. Here we create a list of words and use list comprehension to capitalize each item, split up items with multiple words so each item is a single word, and delete the extraneous *or*:

```
1 x = [word.capitalize()
2     for line in ("hello world?", "world!", "or not")
3     for word in line.split()]
```

```

4     if not word.startswith("or")]
5 >>> x
6 ['Hello', 'World?', 'World!', 'Not']

```

Here two for loops are used: the first one iterates over the text lines, while the second one iterates over words in each of those lines. The final if statement filters out words that start with *or* so they are not included in the final list.

Using list comprehension rather than for loops is a neat way to define lists quickly. Since we're still talking about functional programming, it's worth noting that lists built through list comprehension shouldn't rely on changing the program's state: you are not expected to modify any variable while building the list. This usually makes the lists more concise and easier to read than lists made without listcomp.

Note that there's also syntax for building dictionaries or sets in the same fashion, like so:

```

1 >>> {x:x.upper() for x in ['hello', 'world']}
2 {'world': 'WORLD', 'hello': 'HELLO'}
3 >>> {x.upper() for x in ['hello', 'world']}
4 set(['WORLD', 'HELLO'])

```

8.4 Functional Functions Functioning

You might repeatedly encounter the same set of problems when manipulating data using functional programming. To help you deal with this situation efficiently, Python includes a number of functions for functional programming. This section will give you a quick overview of some of these built-in functions that allows you to build fully functional programs. Once you have an idea of what's available, I encourage you to research further and try out functions where they might apply in your own code.

8.4.1 Applying Functions to Items with `map`

The `map()` function takes the form `map(function, iterable)` and applies function to each item in iterable to return a list in Python 2 or an iterable map object in Python 3, as shown in Example 8.7.

Example 8.7 *map* usage in Python 3

```

1 >>> map(lambda x: x + "bzz!", ["I think", "I'm good"])
2 <map object at 0x7fe7101abdd0>
3 >>> list(map(lambda x: x + "bzz!", ["I think", "I'm good"]))
4 ['I thinkbzz!', "I'm goodbzz!"]

```

You could also write an equivalent of `map()` using list comprehension, which would look like this:

Equivalent of `map` using list comprehension

```

1 >>> (x + "bzz!" for x in ["I think", "I'm good"])
2 <generator object <genexpr> at 0x7f9a0d697dc0>
3 >>> [x + "bzz!" for x in ["I think", "I'm good"]]
4 ['I thinkbzz!', "I'm goodbzz!"]

```

8.4.2 Filtering Lists with `filter`

The `filter()` function takes the form `filter(function or None, iterable)` and filters the items in iterable based on the result returned by function. This will return a list in Python 2 or an iterable filter object in Python 3:

Example 8.8 *filter* usage in Python 3

```

1 >>> filter(lambda x: x.startswith("I "), ["I think", "I'm good"])
2 <filter object at 0x7f9a0d636dd0>
3 >>> list(filter(lambda x: x.startswith("I "), ["I think", "I'm good"])) ←
4 )

```



```
4 ['I think']
```

You could also write an equivalent of `filter()` using list comprehension, like so:

Example 8.9 Equivalent of `filter` using list comprehension

```
1 >>> (x for x in ["I think", "I'm good"] if x.startswith("I "))
2 <generator object <genexpr> at 0x7f9a0d697dc0>
3 >>> [x for x in ["I think", "I'm good"] if x.startswith("I ")]
4 ['I think']
```

8.4.3 Getting Indexes with `enumerate`

The `enumerate()` function takes the form `enumerate(iterable[, start])` and returns an iterable object that provides a sequence of tuples, each consisting of an integer index (starting with `start`, if provided) and the corresponding item in iterable. This function is useful when you need to write code that refers to array indexes. For example, instead of writing this:

```
1 i = 0
2 while i < len(mylist):
3     print("Item %d: %s" % (i, mylist[i]))
4     i += 1
```

you could accomplish the same thing more efficiently with `enumerate()`, like so:

```
1 for i, item in enumerate(mylist):
2     print("Item %d: %s" % (i, item))
```

8.4.4 Sorting a List with `sorted`

The `sorted()` function takes the form `sorted(iterable, key=None, reverse=False)` and returns a sorted version of iterable. The `key` argument allows you to provide a

function that returns the value to sort on:

```
1 >>> sorted([("a", 2), ("c", 1), ("d", 4)])
2 [('a', 2), ('c', 1), ('d', 4)]
3 >>> sorted([("a", 2), ("c", 1), ("d", 4)], key=lambda x: x[1])
4 [('c', 1), ('a', 2), ('d', 4)]
```

8.4.5 Finding Items That Satisfy Conditions with `any` and `all`

The `any(iterable)` and `all(iterable)` functions both return a Boolean depending on the values returned by `iterable`. These simple functions are equivalent to the following full Python code:

```
1 def all(iterable):
2     for x in iterable:
3         if not x:
4             return False
5     return True
```

```
1 def any(iterable):
2     for x in iterable:
3         if x:
4             return True
5     return False
```

These functions are useful for checking whether any or all of the values in an `iterable` satisfy a given condition. For example, the following checks a list for two conditions:

```
1 mylist = [0, 1, 3, -1]
2 if all(map(lambda x: x > 0, mylist)):
3     print("All items are greater than 0")
4 if any(map(lambda x: x > 0, mylist)):
5     print("At least one item is greater than 0")
```

The difference here, as you can see, is that `any()` returns `True` when at least one element meets the condition, while `all()` returns `True` only if every element meets the condition. The `all()` function will also return `True` for an empty iterable, since none of the elements is `False`.

8.4.6 Combining Lists with `zip`

The `zip()` function takes the form `zip(iter1 [,iter2 [...]])` and takes multiple sequences and combines them into tuples. This is useful when you need to combine a list of keys and a list of values into a dict. Like the other functions described here, `zip()` returns a list in Python 2 and an iterable in Python 3. Here we have a list of keys that we map to a list of values to create a dictionary:

```

1 >>> keys = ["foobar", "barzz", "ba!"]
2 >>> map(len, keys)
3 <map object at 0x7fc1686100d0>
4 >>> zip(keys, map(len, keys))
5 <zip object at 0x7fc16860d440>
6 >>> list(zip(keys, map(len, keys)))
7 [('foobar', 6), ('barzz', 5), ('ba!', 3)]
8 >>> dict(zip(keys, map(len, keys)))
9 {'foobar': 6, 'barzz': 5, 'ba!': 3}

```

Warning

You might have noticed by now how the return types differ between Python 2 and Python 3. Most of Python's purely functional built-in functions return a list rather than an iterable in Python 2, making them less memory efficient than their Python 3 equivalents. If you're planning to write code using these functions, keep in mind that you'll get the most benefit out of them in Python 3. If you're stuck to Python 2, don't despair yet: the `itertools` module from the standard library provides an iterator based version of many of these functions (`itertools.izip`, `itertools.imap`, `itertools.ifilter`, and so on).

8.4.7 A Common Problem Solved

There's one important tool still to cover. One common task when working with lists is finding the first item that satisfies a specific condition. There are many ways to accomplish this, and we'll look at those first before I introduce you to the most efficient way: using the `first` package.

8.4.7.1 Finding the Item with Simple Code

Using basic Python code, we might accomplish our aim of finding the first item to satisfy a condition with a function like this:

```
1 def first_positive_number(numbers):
2     for n in numbers:
3         if n > 0:
4             return n
```

In functional style, we could rewrite the `first_positive_number()` function like this:

8.4. FUNCTIONAL FUNCTIONS FUNCTIONING

```
1 def first(predicate, items):
2     for item in items:
3         if predicate(item):
4             return item
5
6 first(lambda x: x > 0, [-1, 0, 1, 2])
```

By using a functional approach where the predicate is passed as argument, the function becomes easily reusable. We could even write it more concisely, like so:

```
1 # Less efficient
2 list(filter(lambda x: x > 0, [-1, 0, 1, 2]))[0] ❶
3 # Efficient but for Python 3
4 next(filter(lambda x: x > 0, [-1, 0, 1, 2]))
5 # Efficient but for Python 2
6 next(itertools.ifilter(lambda x: x > 0, [-1, 0, 1, 2]))
```

❶ Note that this may raise an `IndexError` if no items satisfy the condition, causing `list(filter())` to return an empty list.

For simple cases, you can rely on `next()` to prevent `IndexError` from occurring, like so:

```
1 >>> a = range(10)
2 >>> next(x for x in a if x > 3)
3 4
```

This last example in Example 8.10 will raise `StopIteration` if a condition can never be satisfied. This too can be solved by adding a second argument of `next()`, like so.

Example 8.10 Returning a default value when the condition is not met

```
1 >>> a = range(10)
2 >>> next((x for x in a if x > 10), 'default')
3 'default'
```

Instead of an error, this will return a default value when a condition cannot be met. Luckily for us, Python provides a package to handle all of this for us.

8.4.7.2 Finding the Item Using `first`

Rather than writing out the function from Example 8.10 in every program you make, you can include the small but very useful Python package `first`. This gives you the ability to find the first element of an iterable matching a condition, as shown in Example 8.11.

Example 8.11 Finding the first item in a list that satisfies a condition

```
1 >>> from first import first
2 >>> first([0, False, None, [], (), 42])
3 42
4 >>> first([-1, 0, 1, 2])
5 -1
6 >>> first([-1, 0, 1, 2], key=lambda x: x > 0)
7 1
```

You see that the `first()` function returns the first valid, non-empty item in a list.

8.4.7.3 Using Lambda with `functools`

You'll notice that we've used `lambda` in a good portion of the examples so far in this chapter. The `lambda` function was added to Python to facilitate functional programming functions such as `map()` and `filter()`, which otherwise would have required writing an entirely new function every time you wanted to check a different condition. Example 8.12 is equivalent to Example 8.11 but is written without using `lambda`.

Example 8.12 Using key without lambda

```
1 import operator
2 from first import first
3
4 def greater_than_zero(number):
5     return number > 0
6
7 first([-1, 0, 1, 2], key=greater_than_zero)
```

This code works identically to that in [Example 8.11](#), returning the first non-empty value in a list to meet the condition, but it's a good deal more cumbersome: if we wanted to get the first number in the sequence that's longer than, say, 42 items, we'd need to define an appropriate function via `def` rather than defining it inline with our call to `first`.

But despite its usefulness in helping us avoid situations like this, `lambda` still has its problems. The `first` module contains a `key` argument that can be used to provide a function that receives each item as an argument and returns a Boolean indicating whether it satisfies the condition. However, we can't pass a key function, as it would require more than a single line of code: a `lambda` statement cannot be written on more than one line. That is a significant limitation of `lambda`.

Instead, we would have to go back to the cumbersome pattern of writing new function definitions for each key we need. Or would we?

The `functools` package comes to the rescue with its `partial` method, which provides us with a more flexible alternative to `lambda`. The `functools.partial()` method allows us to create a wrapper function with a twist: rather than changing the behavior of a function, it instead changes the arguments it receives, like so:

```
1 from functools import partial
2 from first import first
3
4 def greater_than(number, min=0):
```

```

5     return number > min
6
7 first([-1, 0, 1, 2], key=partial(greater_than, min=42))

```

Here we create a new `greater_than()` function that works just like the old `greater_than_zero()` from Example 8.12 by default, but this version allows us to specify the value we want to compare our numbers to, whereas before it was hardcoded. Here, we pass `functools.partial()` to our function and the value we want for `min`, and we get back a new function that has `min` set to 42, just as we want. In other words, we can write a function and use `functools.partial()` to customize the behavior of our new functions to suit our needs in any given situation.

Even this version can be pared down. All we're doing in this example is comparing two numbers, and as it turns out, the `operator` module has built-in functions for exactly that:

```

1 import operator
2 from functools import partial
3 from first import first
4
5 first([-1, 0, 1, 2], key=partial(operator.le, 0))

```

This is a good example of `functools.partial()` working with positional arguments. In this case, the function `operator.le(a, b)`, which takes two numbers and returns a Boolean that tells us whether the first number is less than or equal to the second, is passed to `functools.partial()`. The `0` we pass to `functools.partial()` gets assigned to `a`, and the argument passed to the function returned by `functools.partial()` gets assigned to `b`. So this works identically to Example 8.12 but without using `lambda` or defining any additional functions.

**Note**

The `functools.partial` method is typically useful in replacement of `lambda`, and should be considered as a superior alternative. The `lambda` function is something of an anomaly in the Python language, and dropping it altogether was considered for Python 3 due to the function's limited body size of a single line.

8.5 Using `itertools` Functions

Finally, we'll look at some useful functions in the `itertools` module in the Python Standard Library that you should be aware of. Too many programmers end up writing their own versions of these functions simply because they aren't aware that Python provides them out of the box. They are all designed to help you manipulate iterator (that's why the module is called *iter-tools*) and therefore are all purely functional. Here I'll list a few of them and give a brief overview of what they do, and I encourage you to look into them further if they seem of use.

- `accumulate(iterable[, func])` returns a series of accumulation of items from iterables, or whatever is mapped to the `+` operator.
- `chain(*iterables)` iterates over multiple iterables, one after another without building an intermediate list of all items.
- `combinations(iterable, r)` generates all combination of length `r` from the given iterable.
- `compress(data, selectors)` applies a Boolean mask from `selectors` to `data` and returns only the values from `data` where the corresponding element of `selectors` is true.

8.5. USING ITERTOOLS FUNCTIONS

- `count(start, step)` generates an endless sequence of values, starting from `start` and incrementing a `step` at a time with each call.
- `cycle(iterable)` loops repeatedly over the values in `iterable`.
- `repeat(elem[, n])` repeats an element `n` times.
- `dropwhile(predicate, iterable)` filters elements of an iterable starting from the beginning until `predicate` is false.
- `groupby(iterable, keyfunc)` creates an iterator grouping items by the result returned by the `keyfunc` function.
- `permutations(iterable[, r])` returns successive `r`-length permutations of the items in `iterable`.
- `product(*iterables)` returns an iterable of the Cartesian product of iterables without using a nested for loop.
- `takewhile(predicate, iterable)` returns elements of an iterable starting from the beginning until `predicate` is false.

These functions are particularly useful in conjunction with the `operator` module. When used together, `itertools` and `operator` can handle most situations that programmers typically rely on `lambda` for. Here's an example of using `operator.itemgetter()` instead of writing `lambda x:x['foo']`:

Example 8.13 Using the `operator` module with `itertools.groupby`

```
1 >>> import itertools
2 >>> a = [{'foo': 'bar'}, {'foo': 'bar', 'x': 42}, {'foo': 'baz', 'y': ←
    43}]
3 >>> import operator
4 >>> list(itertools.groupby(a, operator.itemgetter('foo')))
5 [('bar', <itertools._grouper object at 0xb000d0>), ('baz', <itertools. ←
    _grouper object at 0xb00110>)]
```

```

6 >>> [(key, list(group)) for key, group in itertools.groupby(a, ←
    operator.itemgetter('foo'))]
7 [ ('bar', [{ 'foo': 'bar' }, { 'x': 42, 'foo': 'bar' }]), ( 'baz', [{ 'y': ←
    43, 'foo': 'baz' }])]

```

In this case, we could have also written `lambda x:x['foo']`, but using `operator` lets us avoid having to use `lambda` at all.

8.6 Summary

While Python is often advertised as being object oriented, it can be used in a very functional manner. A lot of its built-in concepts, such as generators and list comprehension, are functionally oriented and don't conflict with an object-oriented approach. They also limit the reliance on a program's global state, for your own good.

Using functional programming as a paradigm in Python can help you make your program more reusable and easier to test and debug, supporting the *Don't Repeat Yourself* (DRY) mantra. In this spirit, the standard Python modules `itertools` and `operator` are good tools to improve the readability of your functional code.

CHAPTER 9



The Abstract Syntax Tree, Hy, And Lisp-Like Attributes

The abstract syntax tree (AST) is a representation of the structure of the source code of any programming language. Every language, including Python, has a specific AST; Python's AST is built by parsing a Python source file. Like any tree, this one is made of nodes linked together. Each node can contain references to other nodes that make up the tree. A node can represent either an operation, a statement, an expression, or even a module.

This area of Python is not heavily documented and so is not easy to deal with at first glance, but understanding some deeper aspects of how Python is constructed is useful for mastering its usage.

This chapter will examine the AST of some simple Python commands to get you familiar with the structure and how it's used. Once you're familiar with the AST, we'll build a program that can check for wrongly declared methods using flake8 and the AST. Finally, we'll look at Hy, a Python-Lisp hybrid language built on the Python AST.

9.1 Looking at the AST

The easiest way to view the Python AST is to parse some Python code and dump the generated AST. For that, the Python `ast` module provides everything you need, as shown in Example 9.1.

Example 9.1 Parsing Python code to AST

```

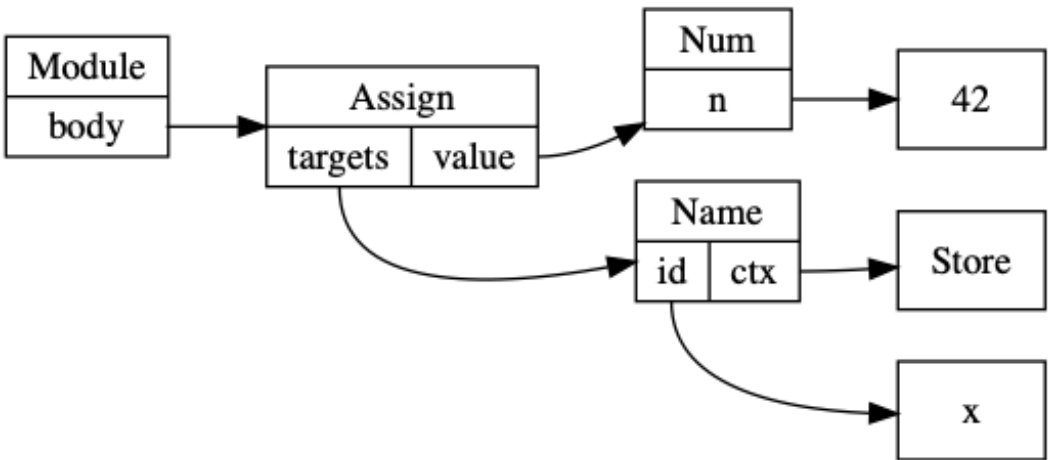
1 >>> import ast
2 >>> ast.parse
3 <function parse at 0x7f062731d950>
4 >>> ast.parse("x = 42")
5 <_ast.Module object at 0x7f0628a5ad10>
6 >>> ast.dump(ast.parse("x = 42"))
7 "Module(body=[Assign(targets=[Name(id='x', ctx=Store())], value=Num(n ↵
    =42))])"
```

The `ast.parse()` function parses any string that contains Python code and returns an `_ast.Module` object. That object is actually the root of the tree: you can browse it to discover every node making up the tree

To visualize what the tree looks like, you can use the `ast.dump()` function, which will return a string representation of the whole tree.

In Example 9.1, the code `x =42` is parsed with `ast.parse()`, and the result is printed using `ast.dump()`. This abstract syntax tree can be rendered as shown in <<example-ast-tree>, which shows the structure of the Python assign command.

Example 9.2 The AST of the assign command in Python



The AST always starts with a root element, which is usually an `_ast.Module` object. This module object contains a list of statements or expressions to evaluate in its `body` attribute and usually represents the content of a file.

As you can probably guess, the `ast.Assign` object shown in Example 9.2 represents an assignment, which is mapped to the `=` sign in the Python syntax. An `ast.Assign` object has a list of targets and a value to set the targets to. The list of targets in this case consists of one object, `ast.Name`, which represents a variable whose ID is `x`. The value is a number `n` with a value (in this case) `42`. The `ctx` attribute stores a context, either `ast.Store` or `ast.Load`, depending on whether the variable is being used for reading or writing. In this case, the variable is being assigned a value, so an `ast.Store` context is used.

This AST would be passed to Python to be compiled and then evaluated, and these tasks can be done via the built-in `compile()` function. This function takes an AST as argument, the source filename, and a mode (either `'exec'`, `'eval'`, or `'single'`). The source filename can be any name that you want your AST to appear to be from; it is common to use the string `'<input>'` as the source filename if the data does not come from a stored file, as shown in Example 9.3.

Example 9.3 Compiling a Python AST

```

1 >>> compile(ast.parse("x = 42"), '<input>', 'exec')
2 <code object <module> at 0x111b3b0, file "<input>", line 1>
3 >>> eval(compile(ast.parse("x = 42"), '<input>', 'exec'))
4 >>> x
5 42

```

The modes stand for execute (`exec`), evaluate (`eval`), and single statement (`single`). The mode should match what has been given to `ast.parse()`, whose default is `exec`.

- The `exec` mode is the normal Python mode, used when an `_ast.Module` is the root of the tree.
- The `eval` mode is a special mode that expects a single `ast.Expression` as the tree.
- Finally, `single` is another special mode that expects a single statement or expression. If it gets an expression, `sys.displayhook()` will be called with the result, as when code is run in the interactive shell.

The root of the AST is `ast.Interactive`, and its `body` attribute is a list of nodes. We could build an AST manually using the classes provided in the `ast` module. Obviously, this is a very long way to write Python code and not a method I would recommend! Nonetheless, it's fun to do and helpful for learning about the AST. Let's see what programming with the AST would look like.

9.2 Writing a Program Using the AST

Let's write a good old "Hello world!" program in Python by building an abstract syntax tree manually.

Example 9.4 Hello world using Python AST

```

1 >>> hello_world = ast.Str(s='hello world!', lineno=1, col_offset=1)
2 >>> print_name = ast.Name(id='print', ctx=ast.Load(), lineno=1, ←
   col_offset=1)
3 >>> print_call = ast.Call(func=print_name, ctx=ast.Load(),
4 ... args=[hello_world], keywords=[], lineno=1, col_offset=1)
5 >>> module = ast.Module(body=[ast.Expr(print_call,
6 ... lineno=1, col_offset=1)], lineno=1, col_offset=1)
7 >>> code = compile(module, '', 'exec')
8 >>> eval(code)
9 hello world!

```

In Example 9.4, we build the tree one leaf at a time, where each leaf is an element (whether a value or an instruction) of the program.

The first leaf is a simple string: the `ast.Str` represents a literal string, which here contains the `hello world!` text.

The `print_name` variable contains an `ast.Name` object, which refers to a variable — in this case, the `print` variable that points to the `print()` function.

The `print_call` variable contains a function call. It refers to the function name to call, the regular arguments to pass to the function call, and the keyword arguments. Which arguments are used depend on the functions being called. In this case, since it's the `print()` function, we'll pass the string we made and stored in `hello_world`.

At last, we create an `ast.Module` object to contain all this code as a list of one expression.

`ast.Module` objects can be compiled using the `compile()` function, which parses the tree and generates a native code object. These code objects are compiled Python code and can finally be executed by a Python virtual machine using `eval`!

This whole process is exactly what happens when you run Python on a `.py` file: once the text tokens are parsed, they are converted into a tree of `ast` objects, compiled, and evaluated.

Note

The arguments `lineno` and `col_offset` represent the line number and column offset, respectively, of the source code that has been used to generate the AST. It doesn't make much sense to set these values in this context since we are not parsing a source file, but it can be useful to be able to find the position of the code that generated the AST. For example, Python uses this information when generating backtraces. Indeed, Python refuses to compile an AST object that doesn't provide this information, so we pass fake values to these. You could also use the `ast.fix_missing_locations()` function to set the missing values to the ones set on the parent node.

9.2.1 The AST Objects

You can view the whole list of objects available in the AST by reading the `_ast` module documentation (note the underscore).

The objects are organized into two main categories: statements and expressions. Statements include types such as `assert`, assignment (`=`), augmented assignment (`+=`, `/=`, etc.), `global`, `def`, `if`, `return`, `for`, `class`, `pass`, `import`, `raise`, and so forth. Statements inherit from `ast.stmt`. Statements influence the control flow of a program and are often composed of expressions.

Expressions include types such as `lambda`, `number`, `yield`, `name` (variable), `comp` are, and `call`. Expressions inherit from `ast.expr`. Expressions differ from statements in that they usually produce a value and have no impact on the program flow.

There are also a few smaller categories, such as the `ast.operator` class, which defines standard operators such as add (`+`), div (`/`), and right shift (`>>`), and the `ast.cmpop` module, which defines comparisons operators.

The simple example here should give you an idea of how to build an AST from scratch. It's easy to then imagine how you might leverage this AST to construct a compiler that would parse strings and generate code, allowing you to implement your own syntax to Python! This is exactly what led to the development of the Hy project, which we'll discuss later in this chapter.

9.2.2 Walking Through an AST

To follow how a tree is built or access particular nodes, you sometimes need to walk through your tree, browsing it and iterating over the nodes. You can do this with the `ast.walk()` function. Alternatively, the `ast` module also provides `NodeTransformer`, a class that you can subclass to walk through an AST and modify particular nodes. Using `NodeTransformer` makes it easy to change code dynamically, as shown in Example 9.5.

Example 9.5 Walking a tree with `NodeTransformer` to alter a node

```

1 import ast
2
3
4 class ReplaceBinOp(ast.NodeTransformer):
5     """Replace operation by addition in binary operation"""
6     def visit_BinOp(self, node):
7         return ast.BinOp(left=node.left,
8                           op=ast.Add(),
9                           right=node.right)
10
11
12 tree = ast.parse("x = 1/3")
13 ast.fix_missing_locations(tree)
14 eval(compile(tree, '', 'exec'))
15 print(ast.dump(tree))
16
17 tree = ReplaceBinOp().visit(tree)

```

9.2. WRITING A PROGRAM USING THE AST

```
18 ast.fix_missing_locations(tree)
19 print(ast.dump(tree))
20 eval(compile(tree, '', 'exec'))
```

The first tree object built is an AST that represents the expression $x = 1/3$. Once this is compiled and evaluated, the result of printing x at the end of the function is `0.33333`, the expected result of $1/3$.

The second tree object is an instance of `ReplaceBinOp`, which inherits from `ast.NodeTransformer`. It implements its own version of the `ast.NodeTransformer.visit()` method and changes any `ast.BinOp` operation to an `ast.BinOp` that executes `ast.Add`. Concretely, this changes any binary operator (+, -, /, and so on) to the + operator. When this second tree is compiled and evaluated, the result is now `4`, which is the result of $1 + 3$, because the / in the first object is replaced with +. You can see the execution of the program here:

```
1 Module(body=[Assign(targets=[Name(id='x', ctx=Store())],
2                     value=BinOp(left=Num(n=1), op=Div(), right=Num(n=3)) ←
3                     )])
4 0.3333333333333333
5 Module(body=[Assign(targets=[Name(id='x', ctx=Store())],
6                     value=BinOp(left=Num(n=1), op=Add(), right=Num(n=3)) ←
7                     )])
8 4
```



Tip

If you need to evaluate a string of Python that should return a simple data type, you can use `ast.literal_eval`. As a safer alternative too `eval`, it prevents the input string from executing any code.

9.3 Extending flake8 with AST Checks

In Section 7.3, you learned that methods that do not rely on the object state should be declared static with the `@staticmethod` decorator. The problem is that a lot of developers simply forget to do so. I've personally spent too much time reviewing code and asking people to fix this problem.

We've seen how to use flake8 to do some automatic checking in the code. It turns out that flake8 is extensible and can provide even more checks. We'll write a flake8 extension that checks for static method declaration omission by analyzing the AST.

Example 9.6 shows an example of one class that omits the static declaration and one that correctly includes it. Write this program out and save it as `ast_ext.py`; we'll use it in a moment to write our extension.

Example 9.6 Omitting and including `@staticmethod`

```

1 class Bad(object):
2     # self is not used, the method does not need
3     # to be bound, it should be declared static
4     def foo(self, a, b, c):
5         return a + b - c
6
7 class OK(object):
8     # This is correct
9     @staticmethod
10    def foo(a, b, c):
11        return a + b - c

```

Though the `Bad.foo` method works fine, strictly speaking it is more correct to write it as `OK.foo` (turn back to Section 7.3 for more detail on why). To check whether all the methods in a Python file are correctly declared, we need to do the following:

- Iterate over all the statement nodes of the AST.

- Check that the statement is a class definition (`ast.ClassDef`).
- Iterate over all the function definitions (`ast.FunctionDef`) of that class statement to check if it is already declared with `@staticmethod`.
- If the method is not declared static, we need to check if the first argument (`self`) is used somewhere in the method. If `self` is not used, the method can be tagged as potentially miswritten.

The name of our project will be `ast_ext`. To register a new plugin in `flake8`, we need to create a packaged project with the usual `setup.py` and `setup.cfg` files. Then, we just need to add an entry point in the `setup.cfg` of our `ast_ext` project.

Example 9.7 Registering new extensions for `flake8`

```
1 [entry_points]
2
3 flake8.extension =
4     [...]
5     H904 = ast_ext:StaticmethodChecker
6     H905 = ast_ext:StaticmethodChecker
```

In Example 9.7, we also register two `flake8` error codes. As you'll notice later, we are actually going to add an extra check to our code while we're at it! The next step is to write the plugin.

9.3.1 Writing the Class

Since we are writing a `flake8` check of the AST, the plugin needs to be a class following a certain signature, as shown in Example 9.8.

Example 9.8 Template for flake8 AST check

```

1 class StaticmethodChecker(object):
2     def __init__(self, tree, filename):
3         self.tree = tree
4
5     def run(self):
6         pass

```

The default template is easy to understand: it stores the tree locally so it can use it in the `run()` method and yield the problems that are discovered. The value that will be yield must follow the expected pep8 signature: a tuple of the form `(lineno, col_offset, error_string, code)`.

9.3.2 Ignoring Irrelevant Code

As indicated earlier, the `ast` module provides the `walk` function, which allows you to iterate easily on a tree. We'll use that to walk through the AST and find out what to check and what not to check.

First, let's write a loop that ignores the statement that are not class definition. Add this to your `ast_ext` project, as shown in [Example 9.8](#).

Example 9.9 Ignoring statements that are not class definitions

```

1 class StaticmethodChecker(object):
2     def __init__(self, tree, filename):
3         self.tree = tree
4
5     def run(self):
6         for stmt in ast.walk(self.tree):
7             # Ignore non-class
8             if not isinstance(stmt, ast.ClassDef):
9                 continue

```

The code in Example 9.9 is still not checking for anything, but now it knows how to ignore statements that are not class definitions. The next step is to set our checker to ignore anything that is not a function definition.

Example 9.10 Ignoring statements that are not function definitions

```

1 for stmt in ast.walk(self.tree):
2     # Ignore non-class
3     if not isinstance(stmt, ast.ClassDef):
4         continue
5     # If it's a class, iterate over its body member to find methods
6     for body_item in stmt.body:
7         # Not a method, skip
8         if not isinstance(body_item, ast.FunctionDef):
9             continue

```

In Example 9.10, we ignore irrelevant statements by iterating over the attributes of the class definition.

9.3.3 Checking for the Correct Decorator

We're all set to write the checking method, which is stored in the `body_item` attribute. First, we need to check whether the method that's being checked is already declared as static. If it is, we don't have to do any further checking and can bail out.

Example 9.11 Checking for the static decorator

```

1 for stmt in ast.walk(self.tree):
2     # Ignore non-class
3     if not isinstance(stmt, ast.ClassDef):
4         continue
5     # If it's a class, iterate over its body member to find methods
6     for body_item in stmt.body:
7         # Not a method, skip
8         if not isinstance(body_item, ast.FunctionDef):

```

```

9         continue
10        # Check that it has a decorator
11        for decorator in body_item.decorator_list:
12            if (isinstance(decorator, ast.Name)
13                and decorator.id == 'staticmethod'):
14                # It's a static function, it's OK
15                break
16        else:
17            # Function is not static, we do nothing for now
18            pass

```

Note that in Example 9.11, we use the special `for/else` form of Python, where the `else` is evaluated unless we use `break` to exit the `for` loop. At this point, we're able to detect whether a method is declared static.

9.3.4 Looking for `self`

The next step is to check whether the method that isn't declared as static uses the `self` argument. First, check whether the method includes any arguments at all, as shown in Example 9.12.

Example 9.12 Checking the method for argumen

```

1  [...]
2  for decorator in body_item.decorator_list:
3      if (isinstance(decorator, ast.Name)
4          and decorator.id == 'staticmethod'):
5          # It's a static function, it's OK
6          break
7  else:
8      try:
9          first_arg = body_item.args.args[0]
10         except IndexError:
11             yield (
12                 body_item.lineno,

```



```

13         body_item.col_offset,
14         "H905: method misses first argument",
15         "H905",
16     )
17     # Check next method
18     continue

```

We finally added a check! This `try` statement in Example 9.12 grabs the first argument from the method signature. If a first argument doesn't exist and so the code fails to retrieve the first argument from the signature, we already know there's a problem: you can't have a bound method without the `self` argument. If the plugin detects that case, it raises the H905 error code we set earlier, signaling a method that misses its first argument.



Note

pep8 codes follow a particular format for error codes (a letter followed by a number), but there are no rules as to which code to pick. You could come up with any other code for this error, as long as it's not already used by pep8 or another extension.

Now you know why we registered two error codes in `setup.cfg`: we had a good opportunity to kill two birds with one stone.

The next step is to check whether the `self` argument is used in the code of the method.

Example 9.13 Checking the method for the self argument

```

1  [...]
2      try:
3          first_arg = body_item.args.args[0]
4      except IndexError:
5          yield (
6              body_item.lineno,

```

9.3. EXTENDING FLAKE8 WITH AST CHECKS

```
7         body_item.col_offset,
8         "H905: method misses first argument",
9         "H905",
10    )
11    # Check next method
12    continue
13    for func_stmt in ast.walk(body_item):
14        if six.PY3:
15            if (isinstance(func_stmt, ast.Name)
16                and first_arg.arg == func_stmt.id):
17                # The first argument is used, it's OK
18                break
19        else:
20            if (func_stmt != first_arg
21                and isinstance(func_stmt, ast.Name)
22                and func_stmt.id == first_arg.id):
23                # The first argument is used, it's OK
24                break
25    else:
26        yield (
27            body_item.lineno,
28            body_item.col_offset,
29            "H904: method should be declared static",
30            "H904",
31        )
```

To check whether the `self` argument is used in the method's body, the plugin in Example 9.13 iterates recursively, using `ast.walk` on the body and looking for the use of the variable named `self`. If the variable isn't found, the program finally yields the H904 error code. Otherwise, nothing happens, and the code is considered sane.

**Note**

As you may have noticed, the code walks over the module AST definition several times. There might be some degree of optimization to browsing the AST in only one pass, but I'm not sure it's worth it considering how the tool is actually used. I'll leave that exercise to you, dear reader.

Knowing the Python AST is not strictly necessary for using Python, but it does give powerful insight into how the language is built and how it works. It thus gives you a better understanding of how the code you write is being used under the hood.

9.4 A Quick Introduction to Hy

Now that you have a good understanding of how Python AST works, you can start dreaming of creating a new syntax for Python. You could parse this new syntax, build an AST out of it, and compile it down to Python code.

This is exactly what *Hy* does. **Hy** is a Lisp dialect that parses a Lisp-like language and converts it to regular Python AST, making it fully compatible with the Python ecosystem. You could compare it to what **Clojure** is to Java. Hy could fill a book by itself, so we will only skim over it in this section. Hy uses the syntax and some features of the Lisp family of languages: it's functionally oriented, provides macros, and is easily extensible.

If you're not already familiar with Lisp — and if not, you should be — the Hy syntax will really look familiar. Once you install Hy (by running `pip install hy`), launching the hy interpreter will give you a standard REPL prompt from which you can start to interact with the interpreter, as shown in Example 9.14.

Example 9.14 Interacting with the Hy interpreter

```

1 % hy
2 hy 0.9.10
3 => (+ 1 2)
4 3

```

For those not familiar with the Lisp syntax, parentheses are used to construct lists. If a list is unquoted, it is evaluated: the first element must be a function, and the rest of the items from the list are passed as arguments. Here the code `(+ 1 2)` is equivalent to Python `1 + 2`.

In Hy, most constructs, such as function definitions, are mapped from Python directly.

Example 9.15 Mapping a function definition from Python

```

1 => (defn hello [name]
2 ... (print "Hello world!")
3 ... (print (% "Nice to meet you %s" name)))
4 => (hello "jd")
5 Hello world!
6 Nice to meet you jd

```

As shown in Example 9.15, internally Hy parses the code provided, converts it to a Python AST, compiles it, and evaluates it. Fortunately, Lisp is an easy tree to parse: each pair of parentheses represents a node of the tree, meaning the conversion is actually easier than for the native Python syntax!

Class definition is supported through the `defclass` construct, which is inspired by the Common Lisp Object System (CLOS).

Example 9.16 Defining a class with `defclass`

```

1 (defclass A [object]
2 [[x 42]

```

```

3 [y (fn [self value]
4   (+ self.x value))]]

```

Example 9.16 defines a class named `A`, which inherits from `object`, with a class attribute `x` whose value is 42; then a method `y` returns the `x` attribute plus a value passed as argument.

What's really wonderful is that you can import any Python library directly into Hy and use it with no penalty. Use the `import()` function to import a module, as shown in Example 9.17, just as you would with regular Python.

Example 9.17 Importing regular Python modules

```

1 => (import uuid)
2 => (uuid.uuid4)
3 UUID('f823a749-a65a-4a62-b853-2687c69d0e1e')
4 => (str (uuid.uuid4))
5 '4efa60f2-23a4-4fc1-8134-00f5c271f809'

```

Hy also has more advanced constructs and macros. In Example 9.18, admire what the `cond()` function can do for you instead of the classic but verbose `if/elif/else`.

Example 9.18 Using `cond` instead of `if/elif/else`

```

1 (cond
2   (> somevar 50)
3   (print "That variable is too big!"))
4 (< somevar 10)
5 (print "That variable is too small!"))
6 (true
7   (print "That variable is jussst right!"))

```

The `cond` macro has the following signature: `(cond [condition_expression return_expression] ...)`. Each condition expression is evaluated, starting with the first: as soon as one of the condition expressions returns a true value, the return expression is evaluated and returned. If no return expression is provided, then the

value of the condition expression is returned. Thus, `cond` is equivalent to an `if/elif` construct, except that it can return the value of the condition expression without having to evaluate it twice or store it in a temporary variable!

Hy allows you to jump into the Lisp world without leaving your comfort zone too far behind you, since you're still writing Python. The `hy2py` tool can even show you what your Hy code would look like once translated into Python. While Hy is not widely used, it is a great tool to show the potential of the Python language. If you're interested in learning more, I suggest you check out the online documentation and join the community.

9.5 Summary

Just like any other programming language, Python source code can be represented using an abstract tree. You'll rarely use the AST directly, but when you understand how it works, it can provide a helpful perspective.

9.6 Interview with Paul Tagliamonte

Paul Tagliamonte created Hy in 2013, and, as a Lisp lover, I joined him in this fabulous adventure. Paul is currently a developer at Sunlight Foundation.

“



How did you learn to use the AST correctly, and do you have any advice for people looking at it?

The AST is extremely underdocumented, so most knowledge comes from reverse engineering generated AST. By writing up simple Python scripts, one can use something similar to `import ast; ast.dump(ast.parse("print foo"))` to generate equivalent AST to help with the task. With a bit of guesswork, and some persistence, it's not untenable to build up a basic understanding this way. At some point, I'll take on the task of documenting my understanding of the AST module, but I find writing code is the best way to learn the AST.

How does Python's AST differ between versions and uses?

Python's AST is not private, but it's not a public interface either. No stability is guaranteed from version to version—in fact, there are some rather annoying differences between Python 2 and 3 and even within different Python 3 releases. In addition, different implementations may interpret the AST differently or even have a unique AST. Nothing says Jython, PyPy, or CPython must deal with the Python AST in the same way.

For instance, CPython can handle slightly out-of-order AST entries (by the `lineno` and `col_offset`), whereas PyPy will throw an assertion error. Though sometimes annoying, the AST is generally sane. It's not impossible to build an AST that works on a vast number of Python instances. With a conditional or two, it's only mildly annoying to create an AST that works on CPython 2.6 through 3.3 and PyPy, making this tool quite handy.

What was your process in creating Hy?

I started on Hy following a conversation about how good it would be to have useful a Lisp that compiles to Python rather than Java's JVM (Clojure). A few short days later, and I had the first version of Hy.

This version resembled a Lisp and even worked like a proper Lisp in some ways, but it was slow. I mean, really slow. It was about an order of magnitude slower than native Python, since the Lisp runtime itself was implemented in Python.

Frustrated, I almost gave up, but then a coworker suggested using the AST to implement the runtime, rather than implementing the runtime in Python. This suggestion was the catalyst for the entire project. I spent my entire holiday break in 2012 hacking on Hy. A week or so later, I had with something that resembled the current Hy codebase.

Just after getting enough of Hy working to implement a basic Flask app, I gave a talk at Boston Python about the project, and the reception was incredibly warm — so warm, in fact, that I start to view Hy as a good way to teach people about Python internals, such as how the REPL works, PEP 302 import hooks, and the Python AST. This was a good introduction to the concept of code that writes code.

I rewrote chunks of the compiler to fix some philosophical issues in the process, leading us to the current iteration of the codebase — which has stood up quite well!

Learning Hy is also a good way to begin understanding how to read Lisp. Users can get comfortable with s-expressions in an environment they know and even use libraries they're already using, easing the transition to other Lisps, such as Common Lisp, Scheme, or Clojure.

How interoperable with Python is Hy?

Hy is amazingly interoperable. So much so that `pdb` can properly debug Hy without you having to make any changes at all. I've written Flask apps, Django apps, and modules of all sorts with Hy. Python can import Python, Hy can import Hy, Hy can import Python, and Python can import Hy. This is what really makes Hy unique; other

Lisp variants like Clojure are purely unidirectional. Clojure can import Java, but Java has one hell of a time importing Clojure.

Hy works by translating Hy code (in s-expressions) into the Python AST almost directly. This compilation step means the generated bytecode is fairly sane stuff, which means Python has a very hard time of even telling the module isn't written in Python at all.

Common Lisp-isms, such as `*earmuffs*` or `using-dashes` are fully supported by translating them into a Python equivalent (in this case, `*earmuffs*` becomes `EARMUFFS`, and `using-dashes` becomes `using_dashes`), which means Python doesn't have a hard time using them at all.

Ensuring that we have really good interoperability is one of our highest priorities, so if you see any bugs — file them!

What are the advantages and disadvantages of choosing Hy?

One advantage of Hy is that it has a full macro system, which Python struggles with. Macros are special functions that alter the code during the compile step. This makes it easy to create new domain-specific languages, which are composed of the base language (in this case, Hy/Python) along with many macros that allow uniquely expressive and succinct code.

As for downsides, Hy, by virtue of being a Lisp written in s-expressions, suffers from the stigma of being hard to learn, read, or maintain. People might be averse to working on projects using Hy for fear of its complexity.

Hy is the Lisp everyone loves to hate. Python folks may not enjoy its syntax, and Lispers may avoid it because Hy uses Python objects directly, meaning the behavior of fundamental objects can sometimes be surprising to the seasoned Lisper.

Hopefully people will look past its syntax and consider exploring parts of Python previously untouched.

”

CHAPTER 10



Performances and Optimizations

“

Premature optimization is the root of all evil.

— Donald Knuth *Structured Programming with go to Statements*

”

Optimizing is rarely the first thing you think about when developing, but there always comes a time when optimizing for better performance will be the right thing to do. That's not to say you should write a program with the idea that it will be slow, but thinking about optimization without first figuring out the right tools to use and doing the proper profiling is a waste of time. Here, I'll show you how to use the right approach to write fast code and where to look when more optimization is needed: many developers focus on the wrong thing when doing optimization, trying to guess where Python might be slower or faster. Rather than speculating, this chapter will help you understand how to profile your application so you will know what part of your program is slowing things down and where the bottlenecks are.

10.1 Data Structures

Most programming problems can be solved in an elegant and simple manner, provided that you use the right data structures—and Python provides many data structures to choose from.

Often, programmers are tempted to code custom data structures, but in reality learning to leverage the existing data structures in Python results in a cleaner and more stable solution.

For example, everybody uses `dict`, but how many times have you seen code trying to access a dictionary by catching the `KeyError` exception:

```
1 def get_fruits(basket, fruit):
2     try:
3         return basket[fruit]
4     except KeyError:
5         return None
```

Or by checking whether the key is present first:

```
1 def get_fruits(basket, fruit):
2     if fruit in basket:
3         return basket[fruit]
```

By using the `get` method already provided by the `dict` class, you can avoid having to catch an exception or checking the key's presence in the first place:

```
1 def get_fruits(basket, fruit):
2     return basket.get(fruit)
```

The method `dict.get()` can also return a default value instead of `None`; just call it with a second argument:

```
1 def get_fruits(basket, fruit):
2     # Return the fruit, or a Banana if the fruit cannot be found.
3     return basket.get(fruit, Banana())
```

Many developers are guilty of using basic Python data structures without being aware of all the methods they provide. This is also true for sets; set data structures have methods that can solve many problems that would otherwise need to be addressed by writing nested for/if blocks. For example, developers often use for/if loops to determine whether an item is in a list, like so:

```
1 def has_invalid_fields(fields):
2     for field in fields:
3         if field not in ['foo', 'bar']:
4             return True
5     return False
```

The loop iterates over each item in the list and checks that all items are either 'foo' or 'bar'. But this can be written more efficiently, removing the need for a loop:

```
1 def has_invalid_fields(fields):
2     return bool(set(fields) - set(['foo', 'bar']))
```

This changes the code to convert the fields to a set, and it gets the rest of the set by subtracting the set(['foo', 'bar']). It then converts the set to a boolean value, which indicates whether any items that aren't foo and bar are left over. By using sets, there is no need to iterate over any list and to check items one by one. A single operation on two sets, done internally by Python, is faster.

Python also has more advanced data structures that can greatly reduce the burden of code maintenance. For example, take a look at Example 10.1.

Example 10.1 Adding an entry in a dictionary of sets

```
1 def add_animal_in_family(species, animal, family):
2     if family not in species:
3         species[family] = set()
4     species[family].add(animal)
5
6 species = {}
7 add_animal_in_family(species, 'cat', 'felidea')
```

10.1. DATA STRUCTURES

Sure, this code is perfectly valid, but how many times will your programs require a variation of Example 10.1? Tens? Hundreds?

Python provides the `collections.defaultdict` structure, which solves the problem in an elegant way:

```
1 import collections
2
3 def add_animal_in_family(species, animal, family):
4     species[family].add(animal)
5
6 species = collections.defaultdict(set)
7 add_animal_in_family(species, 'cat', 'felidea')
```

Each time you try to access a nonexistent item from your dict, the `defaultdict` will use the function that was passed as argument to its constructor to build a new value, instead of raising a `KeyError`. In this case, the `set` function is used to build a new set each time we need it.

The `collections` module offers a few more data structures that can be used to solve other kinds of problems. For example, imagine that you want to count the number of distinct items in an iterable. Let's take a look at the `collections.Counter` method, which provides methods that solve this problem:

```
1 >>> import collections
2 >>> c = collections.Counter("Premature optimization is the root of all ↵
3     evil.")
4 >>> c
5 1
6 >>> c['P'] # Returns the name of occurrence of the letter 'P'
7 4
8 >>> c['e'] # Returns the name of occurrence of the letter 'e'
9 7
10 >>> c.most_common(2) # Returns the 2 most common letters
11 [(' ', 7), ('i', 5)]
```

The `collections.Counter` object works with any iterable that has hashable items, removing the need to write your own counting functions. It can easily count the num-

ber of letters in a string and return the top *n* most common items of an iterable. You might have tried to implement something like this on your own if you were not aware it was already provided by Python's standard library. By learning the various data structures Python offers, you will save time and lessen code maintenance.

With the right data structure, the correct methods, and — obviously — an adequate algorithm, your program should perform well. However, if it is not performing well enough, the best way to get clues about where it might be slow and need some optimization is to profile your code.

10.2 Understanding Behavior Through Profiling

Profiling is a form of dynamic program analysis that allows us to understand how a program behaves. It allows us to determine where there might be bottlenecks and a possible need for optimization. A profile of a program takes the form of a set of statistics that describe how often parts of the program execute and for how long.

Python provides a few tools for profiling your program. One, `cProfile`, is part of the Python standard library and does not require installation. We'll also look at the `dis` module, which can disassemble Python code into smaller parts, making it easier to understand what is happening under the hood.

10.2.1 cProfile

Python has included `cProfile` by default since Python 2.5. To use `cProfile`, call it with your program using the syntax `python -m cProfile <program>`. This will load and enable the `cProfile` module, then run the regular program with instrumentation enabled, as shown in Example [10.2](#).

Example 10.2 Using the cProfile module

```

1 $ python -m cProfile myscript.py
2     343 function calls (342 primitive calls) in 0.000 seconds
3
4 Ordered by: standard name
5
6 ncalls  tottime  percall  cumtime  percall  filename:lineno(function)
7      1    0.000    0.000    0.000    0.000  :0(<_getframe>)
8      1    0.000    0.000    0.000    0.000  :0(<len>)
9     104    0.000    0.000    0.000    0.000  :0(<setattr>)
10      1    0.000    0.000    0.000    0.000  :0(<setprofile>)
11      1    0.000    0.000    0.000    0.000  :0(<startswith>)
12     2/1    0.000    0.000    0.000    0.000  <string>:1(<module>)
13      1    0.000    0.000    0.000    0.000  StringIO.py:30(<module>)
14      1    0.000    0.000    0.000    0.000  StringIO.py:42(StringIO)

```

Example 10.2 shows the output of running a simple script with cProfile. This tells you the number of times each function in the program was called and the time spent on its execution. You can also use the `-s` option to sort by other fields; for example, `-s time` would sort the results by internal time.

We can visualize the information generated by cProfile using a great tool called **KCacheGrind**. This tool was created to deal with programs written in C, but luckily we can use it with Python data by converting the data to a call tree.

The cProfile module has an `-o` option that allows you to save the profiling data, and **pyprof2calltree** can convert data from one format to the other. First, install the converter with the following:

```

1 $ pip install pyprof2calltree

```

Then run the converter as shown in Example 10.3 to both convert the data (`-i` option) and run KCacheGrind with the converted data (`-k` option).

File View Go Settings Help

Open Back Forward Up % Relative Cycle Detection Relative to Parent Shorten Templates Ticks

Flat Profile

Search: (No Grouping)

Incl.	Self	Called	Function	Location
168.31	0.70	(0)	<module> bin/ceilometer-send-counter:21	(unknown)
100.00	1.41	(0)	<_import_>	(unknown)
100.00	0.00	25	<_import_>	(unknown)
92.25	0.00	(0)	<module> /home/jd/Source/oslo.config/oslo/_init_.py:17	(unknown)
92.25	0.00	1	<module> /home/jd/Source/oslo.config/oslo/_init_.py:17	(unknown)
91.55	0.00	1	<module> /usr/lib/python2.7/site-packages/ceilometer/pipeline.py:19	(unknown)
88.73	0.00	(0)	<module> /usr/lib/python2.7/site-packages/ceilometer/pipeline.py:19	(unknown)
60.56	0.00	1	<module> ceilometer/pipeline.py:19	(unknown)
59.86	0.00	(0)	<lambda> /usr/lib/python2.7/site-packages/ceilometer/pipeline.py:19	(unknown)
59.86	0.00	221	<lambda> /usr/lib/python2.7/site-packages/ceilometer/pipeline.py:19	(unknown)
59.86	0.00	221	activate /usr/lib/python2.7/site-packages/ceilometer/pipeline.py:19	(unknown)
59.86	0.00	(0)	subscribe /usr/lib/python2.7/site-packages/ceilometer/pipeline.py:19	(unknown)
59.86	0.00	1	subscribe /usr/lib/python2.7/site-packages/ceilometer/pipeline.py:19	(unknown)
59.15	0.00	(0)	<module> ceilometer/pipeline.py:19	(unknown)
58.45	0.00	(0)	activate /usr/lib/python2.7/site-packages/ceilometer/pipeline.py:19	(unknown)
44.37	0.00	1	<module> ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
42.96	0.70	(0)	<module> ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
38.03	0.00	940	_handle_ns /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
37.32	1.41	(0)	_handle_ns /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
36.62	0.70	(0)	_compile /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
36.62	0.00	1250	_compile /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
35.92	0.00	(0)	compile /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
35.92	0.00	77	compile /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
34.51	0.00	(0)	compile /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
30.28	0.00	457	<map>	(unknown)
29.58	1.41	(0)	<map>	(unknown)
29.58	0.00	933	find_module /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
29.58	0.00	221	fixup_namespace_package /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)
29.58	0.00	1198	realpath /usr/lib/python2.7/site-packages/ceilometer/op.../usr/lib/python2.7/site-packages/ceilometer/op...	(unknown)

prepare_service ceilometer/service.py:80 6.34% <module> ceilometer/sample.py:25 <module> /home/jd/Source/oslo.config/oslo/_init_.py:17 92.25% <module> bin/ceilometer-send-counter:21 168.31% <module> ceilometer/pipeline.py:19 60.56%

Ticks	Count	Callee
92.25	1	<module> /home/jd/Source/oslo.config/oslo/_init_.py:17
60.56	1	<module> ceilometer/pipeline.py:19
6.34	1	prepare_service ceilometer/service.py:80
4.93	1	<module> ceilometer/sample.py:25
2.11	1	<module> /home/jd/Source/oslo.config/oslo/config/cfg.py:261
1.41	1	<module> /usr/lib/python2.7/logging/_init_.py:24

Parts Callee Call Graph All Callee Caller Map Machine Code

pyprof2calltreeEzHmo.log [1] - Total Ticks Cost: 142

Once KCacheGrind opens, it will display information that looks like that in Figure 10.1.

Visualizing the results in this way provides a lot of information that will allow you to determine what part of your program might be consuming too many resources. For instance, you can access the call graph and follow the percentage of time spent in each function, allowing you to determine whether optimization is necessary.

The easiest way to read KCacheGrind is to start with the table on the left of the screen, which lists all the functions and methods executed by your program. You can sort these by execution time, then look at the one that consumes the most CPU time and click on it.

The right panels of KCacheGrind can then show you which other functions have called that function and how many times, as well as which functions are being called by the function. The whole call graph of your program with the execution time of each part is available and made easy to navigate.

This should allow you to better understand which parts are consuming a lot of time and might need optimization. The way to optimize the code is up to you and might depend entirely on what your program is trying to achieve!

While this approach — retrieving information about how your program runs and visualizing it — works well to get a macroscopic view of your program, you need a more microscopic view of some parts of the code to inspect its elements more closely. In such a context, I find it better to rely on the `dis` module to find out what's going on behind the scenes.

10.2.2 Disassembling with the `dis` Module

The `dis` module is a disassembler of Python byte code. Taking code apart can be useful to use to understand what's going on behind each line of the code so you can properly optimize it. For example, Example 10.4 shows the `dis.dis` function, which disassembles whichever function you pass as a parameter and prints the list of bytecode instructions that are run by the function:

Example 10.4 Disassembling Python code

```
1 >>> def x():
2 ...     return 42
3 ...
```

10.2. UNDERSTANDING BEHAVIOR THROUGH PROFILING

```
4 >>> import dis
5 >>> dis.dis(x)
6      2          0 LOAD_CONST          1 (42)
7          3 RETURN_VALUE
```

In Example 10.4, the function `x` is disassembled and its constituents, made of bytecode instructions, are printed. There are only two operations here: loading a constant (`LOAD_CONST`), which is 42, and returning that value (`RETURN_VALUE`). As an example of `dis` in action and how it can be useful, we'll define two functions that do the same thing—concatenate three letters—and disassemble them to see how they do their tasks in different ways:

```
1 abc = ('a', 'b', 'c')
2
3 def concat_a_1():
4     for letter in abc:
5         abc[0] + letter
6
7 def concat_a_2():
8     a = abc[0]
9     for letter in abc:
10        a + letter
```

Both functions appear to do exactly the same thing, but if we disassemble them using `dis.dis`, as shown in Example 10.5, we'll see that the generated bytecode is a bit different.

Example 10.5 Disassembling functions that concatenate strings

```
1 >>> dis.dis(concat_a_1)
2      2          0 SETUP_LOOP          26 (to 29)
3          3 LOAD_GLOBAL              0 (abc)
4          6 GET_ITER
5      >>      7 FOR_ITER              18 (to 28)
6          10 STORE_FAST              0 (letter)
```

10.2. UNDERSTANDING BEHAVIOR THROUGH PROFILING

```
8      3          13 LOAD_GLOBAL          0 (abc)
9          16 LOAD_CONST          1 (0)
10         19 BINARY_SUBSCR
11         20 LOAD_FAST          0 (letter)
12         23 BINARY_ADD
13         24 POP_TOP
14         25 JUMP_ABSOLUTE          7
15     >> 28 POP_BLOCK
16     >> 29 LOAD_CONST          0 (None)
17         32 RETURN_VALUE
18
19 >>> dis.dis(concat_a_2)
20     2          0 LOAD_GLOBAL          0 (abc)
21         3 LOAD_CONST          1 (0)
22         6 BINARY_SUBSCR
23         7 STORE_FAST          0 (a)
24
25     3          10 SETUP_LOOP          22 (to 35)
26         13 LOAD_GLOBAL          0 (abc)
27         16 GET_ITER
28     >> 17 FOR_ITER          14 (to 34)
29         20 STORE_FAST          1 (letter)
30
31     4          23 LOAD_FAST          0 (a)
32         26 LOAD_FAST          1 (letter)
33         29 BINARY_ADD
34         30 POP_TOP
35         31 JUMP_ABSOLUTE          17
36     >> 34 POP_BLOCK
37     >> 35 LOAD_CONST          0 (None)
38         38 RETURN_VALUE
```

In the second function in Example [10.5](#), we store `abc[0]` in a temporary variable before running the loop. This makes the bytecode that's executed inside the loop a little smaller than the bytecode for the first function, as we avoid having to do the `abc[0]` lookup for each iteration. Measured using `timeit`, the second version is 10

percent faster than the first function; it takes a whole microsecond less to execute! Obviously this microsecond is not worth optimizing for unless you call this function billions of times, but this is the kind of insight that the `dis` module can provide.

Whether you rely on “tricks” such as storing the value outside the loop depends on the situation — ultimately, it should be the compiler’s work to optimize this kind of thing. On the other hand, as Python is heavily dynamic, it’s difficult for the compiler to be sure that optimization wouldn’t have negative side effects. In Example 10.5, using `abc[0]` will call `abc.__getitem__`, which could have side effects if it has been overridden by inheritance. Depending on the version of the function you use, the `abc.__getitem__` method will be called once or several times, which might make a difference. Therefore, be careful when writing and optimizing your code!

10.3 Defining Functions Efficiently

One common mistake I have found when reviewing code is definitions of functions within functions. This is inefficient because the function is then redefined repeatedly and needlessly. For example, Example 10.6 shows the `y` function being defined multiple times.

Example 10.6 Disassembling a function defined in a function

```

1  >> import dis
2  >>> def x():
3  ...     return 42
4  ...
5  >>> dis.dis(x)
6      2           0 LOAD_CONST           1 (42)
7           3 RETURN_VALUE
8  >>> def x():
9  ...     def y():
10 ...         return 42
11 ...     return y()
```

10.3. DEFINING FUNCTIONS EFFICIENTLY

```
12 ...
13 >>> dis.dis(x)
14      2          0 LOAD_CONST          1 (<code object y at 0x100ce7e30, ↵
      file "<stdin>", line 2>)
15          3 MAKE_FUNCTION          0
16          6 STORE_FAST             0 (y)
17
18      4          9 LOAD_FAST           0 (y)
19          12 CALL_FUNCTION          0
20          15 RETURN_VALUE
```

Example 10.6 shows the calling of `MAKE_FUNCTION`, `STORE_FAST`, `LOAD_FAST`, and `CALL_FUNCTION`, which requires many more opcodes than those needed to return 42, as seen in Example 10.4.

The only case in which you'd need to define a function within a function is when building a function closure, and this is a perfectly identified use case in Python's opcodes with `LOAD_CLOSURE`, as shown in Example 10.7.

Example 10.7 Disassembling a closure

```
1 >>> def x():
2 ...     a = 42
3 ...     def y():
4 ...         return a
5 ...     return y()
6 ...
7 >>> dis.dis(x)
8      2          0 LOAD_CONST          1 (42)
9          3 STORE_DEREF             0 (a)
10
11      3          6 LOAD_CLOSURE        0 (a)
12          9 BUILD_TUPLE             1
13         12 LOAD_CONST              2 (<code object y at 0x100d139b0, ↵
      file "<stdin>", line 3>)
14         15 MAKE_CLOSURE            0
15         18 STORE_FAST             0 (y)
```

```

16      5          21 LOAD_FAST          0 (y)
17
18      24 CALL_FUNCTION 0
19
27 RETURN_VALUE

```

While you probably won't need to use it every day, disassembling code is a handy tool for when you want a closer look at what happens under the hood.

10.4 Ordered Lists and Bisect

Next, let's look at optimizing lists. One common task when using lists is finding a particular item's position in the list. If the list is unsorted, the worst-case scenario for retrieval has a complexity of $O(n)$, meaning that in the worst case, you'll find your item after iterating over every item of the list.

The usual solution for optimizing this problem is to use a sorted list instead. Sorted lists use a bisecting algorithm for lookup to achieve a retrieve time of $O(\log n)$. The idea is to recursively split the list in half and look on which side, left or right, the item must appear in and so which side should be searched next.

Python provides the `bisect` module, which contains a bisection algorithm, as shown in Example 10.8.

Example 10.8 Using `bisect` to find a needle in a haystack

```

1 >>> farm = sorted(['haystack', 'needle', 'cow', 'pig'])
2 >>> bisect.bisect(farm, 'needle')
3 3
4 >>> bisect.bisect_left(farm, 'needle')
5 2
6 >>> bisect.bisect(farm, 'chicken')
7 0
8 >>> bisect.bisect_left(farm, 'chicken')
9 0

```

```

10 >>> bisect.bisect(farm, 'eggs')
11 1
12 >>> bisect.bisect_left(farm, 'eggs')
13 1

```

As shown in Example 10.8, the `bisect.bisect()` function returns the position an element should be inserted in order to keep the list sorted. Obviously, this only works if the list is properly sorted in the beginning with. Initial sorting allows to us get the theoretical index of an item: `bisect()` does not return whether the item is in the list but where the item should be if it is in the list. Retrieving the item at this index will answer the question about whether the item is in the list.

If you wish to insert the element immediately, the `bisect` module provides the `insort_left` and `insort_right` functions that do exactly that as shown in Example 10.9.

Example 10.9 Inserting an item in a sorted list

```

1 >>> farm
2 ['cow', 'haystack', 'needle', 'pig']
3 >>> bisect.insort(farm, 'eggs')
4 >>> farm
5 ['cow', 'eggs', 'haystack', 'needle', 'pig']
6 >>> bisect.insort(farm, 'turkey')
7 >>> farm
8 ['cow', 'eggs', 'haystack', 'needle', 'pig', 'turkey']

```

Using the `bisect` module, you could also create a special `SortedList` class inheriting from `list` to create a list that is always sorted, as shown in Example 10.10.

Example 10.10 A `SortedList` implementation

```

1 import bisect
2 import unittest
3
4

```



```

5 class SortedList(list):
6     def __init__(self, iterable):
7         super(SortedList, self).__init__(sorted(iterable))
8
9     def insert(self, item):
10        bisect.insort(self, item)
11
12    def extend(self, other):
13        for item in other:
14            self.insert(item)
15
16    @staticmethod
17    def append(o):
18        raise RuntimeError("Cannot append to a sorted list")
19
20    def index(self, value, start=None, stop=None):
21        place = bisect.bisect_left(self[start:stop], value)
22        if start:
23            place += start
24        end = stop or len(self)
25        if place < end and self[place] == value:
26            return place
27        raise ValueError("%s is not in list" % value)
28
29
30 class TestSortedList(unittest.TestCase):
31     def setUp(self):
32         self.mylst = SortedList(
33             ['a', 'c', 'd', 'x', 'f', 'g', 'w']
34         )
35
36     def test_sorted_init(self):
37         self.assertEqual(sorted(['a', 'c', 'd', 'x', 'f', 'g', 'w']),
38                             self.mylst)
39
40     def test_sorted_insert(self):
41         self.mylst.insert('z')

```

```

42 self.assertEqual(['a', 'c', 'd', 'f', 'g', 'w', 'x', 'z'],
43                  self.mylist)
44 self.mylist.insert('b')
45 self.assertEqual(['a', 'b', 'c', 'd', 'f', 'g', 'w', 'x', 'z'],
46                  self.mylist)
47
48 def test_index(self):
49     self.assertEqual(0, self.mylist.index('a'))
50     self.assertEqual(1, self.mylist.index('c'))
51     self.assertEqual(5, self.mylist.index('w'))
52     self.assertEqual(0, self.mylist.index('a', stop=0))
53     self.assertEqual(0, self.mylist.index('a', stop=2))
54     self.assertEqual(0, self.mylist.index('a', stop=20))
55     self.assertRaises(ValueError, self.mylist.index, 'w', stop=3)
56     self.assertRaises(ValueError, self.mylist.index, 'a', start=3)
57     self.assertRaises(ValueError, self.mylist.index, 'a', start ←
        =333)
58
59 def test_extend(self):
60     self.mylist.extend(['b', 'h', 'j', 'c'])
61     self.assertEqual(
62         ['a', 'b', 'c', 'c', 'd', 'f', 'g', 'h', 'j', 'w', 'x'],
63         self.mylist)

```

Using a list class like this is slightly slower when it comes to inserting the item, because the program has to look for the right spot to insert it. However, this class is faster at using the index method than its parent. Obviously, one shouldn't use the `list.append()` method on this class: you can't append an item at the end of the list or it could end up unsorted!

Many Python libraries implement various versions of Example 10.10 and for many more data types, such as binary or red-black tree structures. The **blist** and **bintree** Python packages contain code that can be used for these purposes and are a handy alternative to implementing and debugging your own version. In the next section, we'll see how the native tuple data type provided by Python can be leveraged

to make your Python code a little faster.

10.5 Namedtuple and Slots

Often in programming, you'll need to create simple objects that possess only a few fixed attributes. A simple implementation might be something along these lines:

```
1 class Point(object):
2     def __init__(self, x, y):
3         self.x = x
4         self.y = y
```

This definitely gets the job done. However, there is a downside to this approach. Here we're creating a class that inherits from the object class, so by using this Point class, you are instantiating full objects and allocating a lot of memory.

In Python, regular objects store all of their attributes inside a dictionary, and this dictionary is itself stored in the `__dict__` attribute, as shown in Example 10.11.

Example 10.11 How attributes are stored internally in a Python object

```
1 >>> p = Point(1, 2)
2 >>> p.__dict__
3 {'y': 2, 'x': 1}
4 >>> p.z = 42
5 >>> p.z
6 42
7 >>> p.__dict__
8 {'y': 2, 'x': 1, 'z': 42}
```

For Python, the advantage of using a dict is that it allows you to add as many attributes as you want to an object. The drawback is that using a dictionary to store these attributes is expensive in terms of memory — you need to store the object, the

keys, the value references, and everything else. That makes it slow to create and slow to manipulate, with a high memory cost.

As an example of this unnecessary memory usage, consider the following simple class:

```
1 class Foobar(object):
2     def __init__(self, x):
3         self.x = x
```

This creates a simple Point object with a single attribute named x. Let's check the memory usage of this class using the `memory_profiler`, a nice Python package that allows us to see the memory usage of a program line by line, and a small script that creates 100,000 objects, as shown in Example 10.12.

Example 10.12 Memory profiling with objects

```
1 $ python -m memory_profiler object.py
2 Filename: object.py

3
4 Line #    Mem usage    Increment   Line Contents
5 =====
6      5                                @profile
7      6      9.879 MB      0.000 MB   def main():
8      7      50.289 MB     40.410 MB       f = [ Foobar(42) for i in range ↵
          (100000) ]
```

Example 10.12 demonstrates that creating 100,000 of the objects of the Foobar class would consume 40MB of memory. Although 400 bytes per object might not sound that big, when you are creating thousands of objects, the memory adds up.

There is a way to use objects while avoiding this default behavior of dict: classes in Python can define a `__slots__` attribute that will list only the attributes allowed for instances of this class. Instead of allocating a whole dictionary object to store the object attributes, you can use a list object to store them.

If you go through CPython source code and take a look at the `Objects/typeobject.c` file, it is quite easy to understand what Python does when `__slots__` is set on a class. Example 10.13 is an abbreviated version of the function that handles this:

Example 10.13 Extract of *CPython* `Objects/typeobject.c` code

```

1 static PyObject *
2 type_new(PyTypeObject *metatype, PyObject *args, PyObject *kwargs)
3 {
4     [...]
5     /* Check for a __slots__ sequence variable in dict, and count it ←
6        */
7     slots = _PyDict_GetItemId(dict, &PyId__slots__);
8     nslots = 0;
9     if (slots == NULL) {
10         if (may_add_dict)
11             add_dict++;
12         if (may_add_weak)
13             add_weak++;
14     }
15     else {
16         /* Have slots */
17         /* Make it into a tuple */
18         if (PyUnicode_Check(slots))
19             slots = PyTuple_Pack(1, slots);
20         else
21             slots = PySequence_Tuple(slots);
22         /* Are slots allowed? */
23         nslots = PyTuple_GET_SIZE(slots);
24         if (nslots > 0 && base->tp_itemsize != 0) {
25             PyErr_Format(PyExc_TypeError,
26                          "nonempty __slots__ "
27                          "not supported for subtype of '%s'",
28                          base->tp_name);
29             goto error;
30         }
31         /* Copy slots into a list, mangle names and sort them.

```

```

31         Sorted names are needed for __class__ assignment.
32         Convert them back to tuple at the end.a
33     */
34     newslots = PyList_New(nslots - add_dict - add_weak);
35     if (newslots == NULL)
36         goto error;
37     if (PyList_Sort(newslots) == -1) {
38         Py_DECREF(newslots);
39         goto error;
40     }
41     slots = PyList_AsTuple(newslots);
42     Py_DECREF(newslots);
43     if (slots == NULL)
44         goto error;
45 }
46 /* Allocate the type object */
47 type = (PyTypeObject *)metatype->tp_alloc(metatype, nslots);
48 [...]
49 /* Keep name and slots alive in the extended type object */
50 et = (PyHeapTypeObject *)type;
51 Py_INCREF(name);
52 et->ht_name = name;
53 et->ht_slots = slots;
54 slots = NULL;
55 [...]
56 return (PyObject *)type;

```

As you can see in Example 10.13, Python converts the content of `__slots__` into a tuple and then a list that it builds and sorts before converting it back into a tuple to use and store it in the class. This way, Python can retrieve the values quickly, without having to allocate and use an entire dictionary.

It's easy enough to declare and use such a class. All you need to do is to set the `__slots__` attribute to a list of the attributes that will be defined in the class:

Example 10.14 A class declaration using `__slots__`

```

1 class Foobar(object):
2     __slots__ = ('x',)
3
4     def __init__(self, x):
5         self.x = x

```

We can easily compare the memory usage of the two approaches using the `memory_profiler` Python package, as shown in Example 10.15.

Example 10.15 Memory usage of objects using `__slots__`

```

1 % python -m memory_profiler slots.py
2 Filename: slots.py
3
4 Line #    Mem usage    Increment    Line Contents
5 =====
6      7                                @profile
7      8      9.879 MB      0.000 MB    def main():
8      9     21.609 MB     11.730 MB        f = [ Foobar(42) for i in range ←
          (100000) ]

```

Example 10.15 shows that this time, less than 12MB of memory was needed to create 100,000 objects — or fewer than 120 bytes per object. Thus, by using the `__slots__` attribute of Python classes, we can reduce memory usage, so when we are creating a large number of simple objects, the `__slots__` attribute is an effective and efficient choice. However, this technique shouldn't be used for performing static typing by hardcoding the list of attributes of every class: doing so wouldn't be in the spirit of Python programs.

The drawback here is that the list of attributes is now fixed. No new attribute can be added to the `Foobar` class at runtime.

Due to the fixed nature of the attribute list, it's easy enough to imagine classes where the attributes listed would always have a value and where the fields would always be sorted in some way.

This is exactly what occurs in the `namedtuple` class from the `collections` module. This `namedtuple` class allows us to dynamically create a class that will inherit from the `tuple` class, thus sharing characteristics such as being immutable and having a fixed number of entries.

What `namedtuple` provides is the ability to retrieve tuple elements by referencing a named attribute, rather than having to reference them by index. This makes the tuple easier to access for humans, as shown in [Example 10.16](#).

Example 10.16 Declaring a class using `namedtuple`

```

1 >>> import collections
2 >>> Foobar = collections.namedtuple('Foobar', ['x'])
3 >>> Foobar = collections.namedtuple('Foobar', ['x', 'y'])
4 >>> Foobar(42, 43)
5 Foobar(x=42, y=43)
6 >>> Foobar(42, 43).x
7 42
8 >>> Foobar(42, 43).x = 44
9 Traceback (most recent call last):
10   File "<stdin>", line 1, in <module>
11 AttributeError: can't set attribute
12 >>> Foobar(42, 43).z = 0
13 Traceback (most recent call last):
14   File "<stdin>", line 1, in <module>
15 AttributeError: 'Foobar' object has no attribute 'z'
16 >>> list(Foobar(42, 43))
17 [42, 43]
```

[Example 10.16](#) shows how you can create a simple class with just one line of code and then instantiate it. The list of attributes is fixed, and the existing attributes are immutable. We can't change any attributes of objects of this class or add attributes to them, both because the class inherits from `namedtuple` and because the `__slots__` value is set to an empty tuple, avoiding the creation of the `__dict__`.

Since a class like this would inherit from `tuple`, we can easily convert it to a list,

as shown in Example 10.16.

Example 10.17 Memory usage of a class built from `collections.namedtuple`

```

1 % python -m memory_profiler namedtuple.py
2 Filename: namedtuple.py
3
4 Line #    Mem usage    Increment    Line Contents
5 =====
6         4                                @profile
7         5      9.895 MB      0.000 MB    def main():
8         6     23.184 MB     13.289 MB      f = [ Foobar(42) for i in range ←
          (100000) ]

```

Example 10.17 demonstrates the memory usage of the `namedtuple` class factory. At around 13MB for 100,000 objects, using `namedtuple` is slightly less efficient than using an object with `__slots__`, but the bonus is that it is compatible with the `tuple` class. It can therefore be passed to many native Python functions and libraries that expect an iterable as an argument. A `namedtuple` class factory also enjoys the various optimizations that exist for tuples: for example, tuples with fewer items than `PyTuple_MAXSAVESIZE` (20 by default) will use a faster memory allocator in CPython.

The `namedtuple` class also provides a few extra methods that, even if prefixed by an underscore, are actually intended to be public. The `_asdict()` method can convert the `namedtuple` to a `dict` instance, the `_make()` method allows you to convert an existing iterable object to this class, and `_replace()` returns a new instance of the object with some fields replaced.

Namedtuples are a great replacement for small objects that consists of only a few attributes and do not necessitate any custom methods — think about them rather than using dictionaries, for example. If your data type needs methods, has a fixed list of attribute, and might be instantiated thousands of times, then creating a custom class using `__slots__` might be a good idea to save some memory.

10.6 Memoization

Memoization is an optimization technique used to speed up function calls by caching their results. The results of a function can be cached only if the function is pure, meaning that it has no side effects and does not depend on any global state. (See Chapter 8 for more on pure functions.)

One trivial function that can be memoized is `sin()`, shown in Example 10.18.

Example 10.18 A basic memoization

```

1 >>> import math
2 >>> _SIN_MEMOIZED_VALUES = {}
3 >>> def memoized_sin(x):
4 ...     if x not in _SIN_MEMOIZED_VALUES:
5 ...         _SIN_MEMOIZED_VALUES[x] = math.sin(x)
6 ...     return _SIN_MEMOIZED_VALUES[x]
7 >>> memoized_sin(1)
8 0.8414709848078965
9 >>> _SIN_MEMOIZED_VALUES
10 {1: 0.8414709848078965}
11 >>> memoized_sin(2)
12 0.9092974268256817
13 >>> memoized_sin(2)
14 0.9092974268256817
15 >>> _SIN_MEMOIZED_VALUES
16 {1: 0.8414709848078965, 2: 0.9092974268256817}
17 >>> memoized_sin(1)
18 0.8414709848078965
19 >>> _SIN_MEMOIZED_VALUES
20 {1: 0.8414709848078965, 2: 0.9092974268256817}

```

In Example 10.18, the first time that `memoized_sin()` is called with an argument that is not stored in `_SIN_MEMOIZED_VALUES`, the value is computed and stored in this dictionary. If we call the function with the same value again, the result will be retrieved from the dictionary rather than recomputed. While `sin()` computes very

quickly, some advanced functions involving more complicated computations may take longer, and this is where memorization really shines.

If you’ve already read about decorators (if not, see Section 7.1), you might see a perfect opportunity to use them here, and you’d be right. PyPI lists a few implementations of memoization through decorators, from very simple cases to the most complex and complete.

Starting with Python 3.3, the `functools` module provides a *least recently used* (LRU) cache decorator. This provides the same functionality as the memoization described earlier, but with the benefit that it limits the number of entries in the cache, removing the least recently used one when the cache reaches its maximum size. The module also provides statistics on cache hits and misses (whether something was in the accessed cache or not), among other data. In my opinion, these statistics are must-haves when implementing such a cache. The strength of using memorization, or any caching technique, is in the ability to meter its usage and usefulness.

Example 10.19 demonstrates how to use the `functools.lru_cache()` method to implement the memoization of a function. When decorated, the function gets a `cache_info()` method that can be called to get statistics about the cache usage.

Example 10.19 Using `functools.lru_cache`

```

1 >>> import functools
2 >>> import math
3 >>> @functools.lru_cache(maxsize=2)
4 ... def memoized_sin(x):
5 ...     return math.sin(x)
6 ...
7 >>> memoized_sin(2)
8 0.9092974268256817
9 >>> memoized_sin.cache_info()
10 CacheInfo(hits=0, misses=1, maxsize=2, currsize=1)
11 >>> memoized_sin(2)
12 0.9092974268256817
13 >>> memoized_sin.cache_info()

```

```

14 CacheInfo(hits=1, misses=1, maxsize=2, currsize=1)
15 >>> memoized_sin(3)
16 0.1411200080598672
17 >>> memoized_sin.cache_info()
18 CacheInfo(hits=1, misses=2, maxsize=2, currsize=2)
19 >>> memoized_sin(4)
20 -0.7568024953079282
21 >>> memoized_sin.cache_info()
22 CacheInfo(hits=1, misses=3, maxsize=2, currsize=2)
23 >>> memoized_sin(3)
24 0.1411200080598672
25 >>> memoized_sin.cache_info()
26 CacheInfo(hits=2, misses=3, maxsize=2, currsize=2)
27 >>> memoized_sin.cache_clear()
28 >>> memoized_sin.cache_info()
29 CacheInfo(hits=0, misses=0, maxsize=2, currsize=0)

```

Example 10.19 demonstrates how your cache is being used and how to tell whether there are optimizations to be made. For example, if the number of misses is high when the cache is not full, then the cache may be useless because the arguments passed to the function are never identical. This will help determine what should or should not be memoized!

10.7 Faster Python with PyPy

PyPy is an efficient implementation of the Python language which complies with standards: you should be able to run any Python program with it. Indeed, the canonical implementation of Python, *CPython* — so called because it's written in C — can be very slow. The idea behind PyPy was to write a Python interpreter in Python itself. In time, it evolved to be written in RPython, which is a restricted subset of the Python language.

RPython places constraints on the Python language in such a way that a vari-

able's type can be inferred at compile time. The RPython code is translated into C code, which is compiled to build the interpreter. RPython could of course be used to implement languages other than Python.

What's interesting in PyPy, besides the technical challenge, is that it is now at a stage where it can act as a faster replacement for CPython. PyPy has a *just-in-time* (JIT) compiler built-in; in other words, it allows the code to run faster by combining the speed of compiled code with the flexibility of interpretation.

How fast? That depends, but for pure algorithmic code, it is much faster. For more general code, PyPy claims to achieve three times the speed of CPython most of the time. Unfortunately, PyPy also has some of the limitations of CPython, including the global interpreter lock (GIL), which allows only one thread to execute at a time.

Though it's not strictly an optimization technique, targeting PyPy as one of your supported Python implementations might be a good idea. To make PyPy a support implementation, you basically need to make sure that you are testing your software under PyPy as you would under CPython. In Section 6.3, we discussed tox which supports the building of virtual environments using PyPy, just as it does for any version of CPython, so putting PyPy support in place should be pretty straightforward.

Testing PyPy support right at the beginning of the project will make sure that there's not too much work to do at a later stage if you decide then that you want to be able to run your software with PyPy.

Note



For the **Hy** project discussed in Section 9.4, we successfully adopted this strategy from the beginning. Hy always has supported PyPy and all other CPython versions without much trouble. On the other hand, the OpenStack project failed to do so for its projects and, as a result, is now blocked by various code paths and dependencies that don't work on PyPy for various reasons; they weren't required to be fully tested in the early stages.

PyPy is compatible with Python 2.7 and Python 3.5, and its JIT compiler works on 32- and 64-bit, x86, and ARM architectures and under various operating systems (Linux, Windows, and macOS). PyPy often lags behind CPython in features, but it regularly catches up. Unless your project is reliant on the latest CPython features, this lag might not be a problem.

10.8 Achieving Zero Copy with the Buffer Protocol

Often programs have to deal with huge amounts of data in the form of large arrays of bytes. Handling such a large quantity of input in strings can be very ineffective once you start manipulating the data by copying, slicing, and modifying it.

Let's consider a small program that reads a large file of binary data and copies it partially into another file. To examine the memory usage of this program, we will use `memory_profiler`, as we did earlier. The script to partially copy the file is shown in Example 10.20.

Example 10.20 Writing a file with random data

```
1 @profile
2 def read_random():
3     with open("/dev/urandom", "rb") as source:
4         content = source.read(1024 * 10000)
5         content_to_write = content[1024:]
6         print("Content length: %d, content to write length %d" %
7               (len(content), len(content_to_write)))
8         with open("/dev/null", "wb") as target:
9             target.write(content_to_write)
10
11 if __name__ == '__main__':
12     read_random()
```

Running the program in Example 10.20 using `memory_profiler` produces the output shown in Example 10.21.

Example 10.21 Running memory profiler

```

1 $ python -m memory_profiler memoryview/copy.py
2 Content length: 10240000, content to write length 10238976
3 Filename: memoryview/copy.py
4
5 Mem usage    Increment   Line Contents
6 =====
7
8                @profile
9  9.883 MB     0.000 MB   def read_random():
10  9.887 MB     0.004 MB       with open("/dev/urandom", "rb") as source:
11 19.656 MB     9.770 MB           content = source.read(1024 * 10000) ❶
12 29.422 MB     9.766 MB           content_to_write = content[1024:] ❷
13 29.422 MB     0.000 MB           print("Content length: %d, content to write ←
14             length %d" %
15             (len(content), len(content_to_write)))
16 29.434 MB     0.012 MB       with open("/dev/null", "wb") as target:
17 29.434 MB     0.000 MB           target.write(content_to_write)

```

- ❶ Reads 10 MB from `_/dev/urandom`. Python needs to allocate around 10 MB of memory to store this data as a string.
- ❷ Copy the entire block of data minus the first KB.

What's interesting in Example 10.21 is that the program's memory usage is increased by about 10MB when building the variable `content_to_write`. In fact, the slice operator is copying the entirety of `content`, minus the first KB, into a new string object, allocating a large chunk of the 10MB.

Performing this kind of operation on large byte arrays is going to be a disaster since large pieces of memory will be allocated and copied. If you have experience writing in C code, you know that using the `memcpy()` function has a significant cost in terms of both memory usage and general performance: copying memory is slow.

But as a C programmer, you'll also know that strings are arrays of characters and that nothing stops you from looking at only part of an array without copying it. You can do this through the use of basic pointer arithmetic—assuming that the entire string is in a contiguous memory area.

This is also possible in Python using objects that implement the buffer protocol. The buffer protocol is defined in [PEP 3118](#), as a C API that needs to be implemented on various types for them to provide this protocol. The string class, for example, implements this protocol.

When you implement this protocol on an object, you can then use the `memoryview` class constructor to build a new `memoryview` object that will reference the original object memory. For example, [Example 10.22](#) shows how to use `memoryview` to access slice of a string without doing any copying:


Example 10.22 Using `memoryview`

```

1 >>> s = b"abcdefgh"
2 >>> view = memoryview(s)
3 >>> view[1]
4 98 ❶
5 >>> limited = view[1:3]
6 >>> limited
7 <memory at 0x7fca18b8d460>
8 >>> bytes(view[1:3])
9 b'bc'
```

❶ This is the ASCII code for the letter *b*.

In [Example 10.22](#), we are making use of the fact that the `memoryview` object's slice operator itself returns a `memoryview` object. That means it does not copy any data but merely references a particular slice of it, saving the memory that would be consumed by a copy. [Figure 10.2](#) illustrates what happens in [Example 10.22](#).

Figure 10.2: Using slice on *memoryview* objects

Knowing that, we can rewrite the program from Example 10.20, this time referencing the data we want to write using a *memoryview* object rather than allocating a new string.

Example 10.23 Writing random data to a file, using less memory

```

1 @profile
2 def read_random():
3     with open("/dev/urandom", "rb") as source:
4         content = source.read(1024 * 10000)
5         content_to_write = memoryview(content)[1024:]
6     print("Content length: %d, content to write length %d" %
7           (len(content), len(content_to_write)))
8     with open("/dev/null", "wb") as target:
9         target.write(content_to_write)
10
11 if __name__ == '__main__':
12     read_random()
  
```

The program in Example 10.23 uses half the memory usage of the first version in Example 10.24. We can see this by testing it with *memory_profiler* again, like so:

Example 10.24 Running memory profiler with the optimized version

```

1 $ python -m memory_profiler memoryview/copy-memoryview.py
2 Content length: 10240000, content to write length 10238976
  
```

10.8. ACHIEVING ZERO COPY WITH THE BUFFER PROTOCOL

```
3 Filename: memoryview/copy-memoryview.py
4
5 Mem usage    Increment  Line Contents
6 =====
7
8                @profile
9 9.887 MB     0.000 MB def read_random():
10 9.891 MB     0.004 MB     with open("/dev/urandom", "rb") as source:
11 19.660 MB    9.770 MB         content = source.read(1024 * 10000) ❶
12 19.660 MB    0.000 MB         content_to_write = memoryview(content) ←
13     [1024:] ❷
14 19.660 MB    0.000 MB         print("Content length: %d, content to write ←
15     length %d" %
16         (len(content), len(content_to_write)))
17 19.672 MB    0.012 MB     with open("/dev/null", "wb") as target:
18 19.672 MB    0.000 MB         target.write(content_to_write)
```

- ❶ We are reading 10 MB from */dev/urandom* and not doing much with it. Python needs to allocate around 10 MB of memory to store this data as a string.
- ❷ We reference the entire block of data minus the first KB – because we won't be writing to that first KB to the target file. No copying means that no more memory is used!

This kind of trick is especially useful when dealing with sockets. When sending data over a socket, it's possible that the data might split between calls rather than be sent in a single call: the `socket.send` methods return the actual data length that was able to be sent by the network, which might be smaller than the data that was aimed to be sent. Example 10.25 shows how the situation is usually handled.

Example 10.25 Reading data from a socket

```
1 import socket
2 s = socket.socket(...)
3 s.connect(...)
4 data = b"a" * (1024 * 10000) ❶
```

10.8. ACHIEVING ZERO COPY WITH THE BUFFER PROTOCOL

```
5 while data:
6     sent = s.send(data)
7     data = data[sent:] ❷
```

- ❶ Build a bytes object with more than 100 millions times the letter *a*.
- ❷ Remove the first *sent* bytes sent.

Using a mechanism as implemented in Example 10.25, a program will copy the data over and over until the socket has sent everything.

We can alter the program in Example 10.25 to use `memoryview` to achieve the same functionality with zero copying, and therefore higher performance as shown in Example 10.26.

Example 10.26 Reading data from a socket using `memoryview`

```
1 import socket
2 s = socket.socket(...)
3 s.connect(...)
4 data = b"a" * (1024 * 100000) ❶
5 mv = memoryview(data)
6 while mv:
7     sent = s.send(mv)
8     mv = mv[sent:] ❷
```

- ❶ Build a bytes object with more than 100 millions times the letter *a*.
- ❷ Build a new `memoryview` object pointing to the data which remains to be sent.

As this program won't copy anything, it won't use any more memory than the 100MB initially needed for the data variable.

We've now seen how `memoryview` objects can be used to write data efficiently, and this same method can be used to read data. Most I/O operations in Python know

10.8. ACHIEVING ZERO COPY WITH THE BUFFER PROTOCOL

how to deal with objects implementing the buffer protocol: they can read from those, but also write to those. In this case, we don't need `memoryview` objects; we can just ask an I/O function to write into our preallocated object, as shown in Example 10.27.

Example 10.27 Reading data from a file into a bytearray

```
1 >>> ba = bytearray(8)
2 >>> ba
3 bytearray(b'\x00\x00\x00\x00\x00\x00\x00\x00')
4 >>> with open("/dev/urandom", "rb") as source:
5 ...     source.readinto(ba)
6 ...
7 8
8 >>> ba
9 bytearray(b'\x00\x00\x00\x00\x00\x00\x00\x00')
```

In Example 10.27, by using the `readinto()` method of the opened file, Python can directly read the data from the file and write it to a preallocated `bytearray`. With such techniques, it's easy to preallocate a buffer (as you would do in C to mitigate the number of calls to `malloc()`) and fill it at your convenience: using `memoryview`, you can place data at any point in the memory area, as shown in Example 10.28.

Example 10.28 Reading data from a file into a bytearray at a specified location

```
1 >>> ba = bytearray(8)
2 >>> ba_at_4 = memoryview(ba)[4:] ❶
3 >>> with open("/dev/urandom", "rb") as source:
4 ...     source.readinto(ba_at_4) ❷
5 ...
6 4
7 >>> ba
8 bytearray(b'\x00\x00\x00\x00\x0b\x19\xae\xb2')
```

❶ We reference the *bytearray* from offset 4 to its end.

- 2 We write the content of `/dev/urandom` from offset 4 to the end of the *bytearray*, effectively reading 4 bytes only.

The buffer protocol is extremely important for achieving low memory overhead and great performances. As Python hides all the memory allocations, developers tend to forget what happens under the hood, at great cost to the speed of their programs!

Both the objects in the `array` module and the functions in the `struct` module can handle the buffer protocol correctly and can therefore perform efficiently when targeting zero copying.

10.9 Summary

As we've seen in this chapter, there are plenty of ways to make Python code faster. Choosing the right data structure and using the correct methods for manipulating the data can have a huge impact in terms of CPU and memory usage. That's why it's very important to understand what happens in Python internally.

However, optimization should never be done prematurely, that is, without first performing a proper profiling. It is too easy to waste time rewriting some barely used code with a faster variant while missing central pain points. Don't miss the big picture.

10.10 Victor Stinner on Optimization

Victor Stinner is a longtime Python hacker, a core contributor, and the author of many Python modules. He authored [PEP 454](#) in 2013, which proposed a new `tracemalloc` module to trace memory block allocation inside Python, and he wrote a

simple AST optimizer called FAT. He also regularly contributes to the improvement of CPython performance.

“

**What’s a good starting strategy for optimizing Python code?**

The strategy is the same in Python as in other languages. First, you need a well-defined use case in order to get a stable and reproducible benchmark. Without a reliable benchmark, trying different optimizations may result in wasted time and premature optimization. Useless optimizations may make the code worse, less readable, or even slower. A useful optimization must speed the program up by at least 5 percent if it’s to be worth pursuing.

If a specific part of the code is identified as being “slow,” a benchmark should be prepared on this code. A benchmark on a short function is usually called a *micro-benchmark*. The speedup should be at least 20 percent, maybe 25 percent, to justify an optimization on a micro-benchmark.

It may be interesting to run a benchmark on different computers, different operating systems, or different compilers. For example, performances of `realloc()` may vary between Linux and Windows.

What are your recommended tools for profiling or optimizing Python code?

Python 3.3 has a `time.perf_counter()` function to measure elapsed time for a benchmark. It has the best resolution available.

A test should be run more than once; three times is a minimum, and five may be enough. Repeating a test fills disk cache and CPU caches. I prefer to keep the minimum timing; other developers prefer the geometric mean.

For micro-benchmarks, the `timeit` module is easy to use and gives results quickly, but the results are not reliable using default parameters. Tests should be repeated manually to get stable results.

Optimizing can take a lot of time, so it's better to focus on functions that use the most CPU power. To find these functions, Python has `cProfile` and `profile` modules to record the amount of time spent in each function.

Do you have any Python tricks that could improve performance?

You should reuse the standard library as much as possible — it's well tested and also usually efficient. Built-in Python types are implemented in C and have good performance. Use the correct container to get the best performance; Python provides many different kind of containers: `dict`, `list`, `deque`, `set`, and so on.

There are some hacks for optimizing Python, but you should avoid these because they make the code less readable in exchange for a minor speedup.

The Zen of Python (PEP 20) says, “There should be one—and preferably only one—obvious way to do it.” In practice, there are different ways to write Python code, and performances are not the same. Only trust benchmarks on your use case.

Which areas of Python have the poorest performance and

should be watched out for?

In general, I prefer not to worry about performance while developing a new application. Premature optimization is the root of all evil. When you identify slow functions, change the algorithm. If the algorithm and the container types are well chosen, you might rewrite short functions in C to get the best performance.

One bottleneck in CPython is the global interpreter lock, known as the GIL. Two threads cannot execute Python bytecode at the same time. However, this limitation only matters if two threads are executing pure Python code. If most processing time is spent in function calls, and these functions release the GIL, then the GIL is not the bottleneck. For example, most I/O functions release the GIL.

The multiprocessing module can easily be used to work around the GIL. Another option, more complex to implement, is to write asynchronous code. Twisted, Tornado, and Tulip projects, which are network-oriented libraries, make use of this technique.

What are some often-seen performance mistakes?

When Python is not well understood, inefficient code can be written. For example, I have seen `copy.deepcopy()` misused, when no copying was required. Another performance killer is an inefficient data structure. With less than 100 items, the container type has no impact on performance. With more items, the complexity of each operation (add, get, delete) and its effects must be known.

”

CHAPTER 11



Scaling and Architecture

Sooner or later, your development process is going to have to consider resiliency and scalability. The scalability, concurrency, and parallelism of an application depend largely on the choices made about its initial architecture and design. As we'll see in this chapter, there are some paradigms—such as multithreading—that don't apply correctly to Python, whereas other techniques, such as service-oriented architecture, work better.

Covering scalability in its entirety would take an entire book, and in fact has been covered by many books, including my follow-up book to this one, *Scaling Python*, dedicated to writing distributed and scalable Python applications. This chapter, therefore, covers just the fundamentals of scaling that you should know, even if you're not planning to build applications with millions of users.

11.1 Multithreading in Python and Its Limitations

By default, Python processes run on one and only one thread, called the main thread. This thread is responsible for executing code on a single processor. Multithreading is a programming technique that allows code to run concurrently inside a single Python

process by running several threads at the same time. This is the primary mechanism through which we can introduce concurrency in Python. If the computer is equipped with multiple processors, you can even use parallelism, running threads in parallel over several processors, to make code execution faster.

Multithreading is most commonly used (though not always appropriately) when:

- You need to run background or I/O-oriented tasks without stopping your main thread's execution. For example, the main loop of a graphical user interface is busy waiting for an event (e.g., a user click or keyboard input) but the code needs to execute other tasks.
- You need to spread your workload across several CPUs.

The first scenario is a good general case for multithreading. Though implementing multithreading in this circumstance would introduce extra complexity, controlling multithreading would be manageable, and there would be little performance penalty as long as the workload was not CPU intensive. The performance gain from using concurrency with workloads that are I/O intensive gets more interesting when the I/O has high latency: the more often you have to wait to read or write, the more beneficial it is to do something else in the meantime.

In the second scenario, you might want to start a new thread for each new request instead of handling them one at a time. This may seem like a good use for multithreading. However, if you try to spread your workload out like this, you will encounter the Python global interpreter lock (GIL), a lock that must be acquired each time CPython needs to execute bytecode. The lock means that only one thread can have control of the Python interpreter at any one time. This rule was introduced originally to prevent race conditions, but it unfortunately means that if you try to scale your application by making it run multiple threads, you'll always be limited by this global lock.

So, while using threads seems like the ideal solution, most applications running requests in multiple threads struggle to attain 150 percent CPU usage, or usage of

the equivalent of 1.5 cores. Most computers have 4 or 8 cores, and servers offer 24 or 48 cores, but the GIL prevents Python from using the full CPU. There are some initiatives underway to remove the GIL, but the effort is extremely complex as it requires trade-offs between performance and backward compatibility.

Although CPython is the most commonly used implementation of the Python language, there are others that do not have a GIL. **Jython**, for example, can efficiently run multiple threads in parallel. Unfortunately, projects such as Jython by their very nature lag behind CPython and so are not really useful targets; innovation happens in CPython, and the other implementations are just following in CPython's footsteps.

So, let's revisit our two use cases with all that we now know and figure out a better solution:

- When you need to run background tasks, you can use multithreading, but the easier solution is to build your application around an event loop. There are a lot of Python modules that provide for this, and the standard is now `asyncio`. There are also frameworks, such as Twisted, built around the same concept. The most advanced frameworks will give you access to events based on signals, timers, and file descriptor activity — we'll talk about this later in the Section [11.3](#).
- When you need to spread the workload, using multiple processes is the most efficient method. We'll look at this technique in the next section.

Developers should always think twice before using multithreading. As one example, I once used multithreading to dispatch jobs in *rebuild*, a Debian build daemon I wrote a few years ago. While it seemed handy to have a different thread to control each running build job, I very quickly fell into the threading-parallelism trap in Python. If I had the chance to begin again, I'd build something based on asynchronous events handling or multiprocessing and not have to worry about the GIL.

Multithreading is complex, and it's hard to get multithreaded applications right. You need to handle thread synchronization and locking, which means there are a lot

of opportunities to introduce bugs. Considering the small overall gain, it's better to think twice before spending too much effort on it.

11.2 Multiprocessing vs. Multithreading

Since the GIL prevents multithreading from being a good scalability solution, look to the alternative solution offered by the multiprocessing package provided with Python. The package exposes the same kind of interface you'd achieve using the multithreading module, except that it starts new processes (via `os.fork()`) instead of new system threads.

[Using multithreading for concurrent activity](#) shows a simple example in which one million random integers are summed eight times, with this activity spread across eight threads at the same time.

Using multithreading for concurrent activity

```

1 import random
2 import threading
3
4 results = []
5
6
7 def compute():
8     results.append(sum(
9         [random.randint(1, 100) for i in range(1000000)]))
10
11
12 workers = [threading.Thread(target=compute) for x in range(8)]
13 for worker in workers:
14     worker.start()
15 for worker in workers:
16     worker.join()
17 print("Results: %s" % results)

```

11.2. MULTIPROCESSING VS. MULTITHREADING

In [Using multithreading for concurrent activity](#), we create eight threads using the `threading.Thread` class and store them into the `workers` array. Those threads will execute the `compute()` function. They then use the `start()` method to start. The `join()` method only returns once the thread has terminated its execution. At this stage, the result can be printed.

Running this program returns the following:

Example 11.1 Result of *time python worker.py*

```
1 $ time python worker.py
2 Results: [50517927, 50496846, 50494093, 50503078, 50512047, 50482863, ↵
   50543387, 50511493]
3 python worker.py 13.04s user 2.11s system 129% cpu 11.662 total
```

This has been run on an idle four-core CPU, which means that Python could potentially have used up to 400 percent of CPU. However, these results show that it was clearly unable to do that, even with eight threads running in parallel. Instead, its CPU usage maxed out at 129 percent, which is just 32 percent of the hardware's capabilities (129/400).

Now, let's rewrite this implementation using multiprocessing. For a simple case like this, switching to multiprocessing is pretty straightforward, as shown in [Example 11.2](#).

Example 11.2 Using multiprocessing for concurrent activity

```
1 import multiprocessing
2 import random
3
4
5 def compute(n):
6     return sum(
7         [random.randint(1, 100) for i in range(1000000)]
8     )
9
```

```

10 # Start 8 workers
11 pool = multiprocessing.Pool(processes=8)
12 print("Results: %s" % pool.map(compute, range(8)))

```

The multiprocessing module offers a `Pool` object that accepts as an argument the number of processes to start. Its `map()` method works in the same way as the native `map()` method, except that a different Python process will be responsible for the execution of the `compute()` function.

Running the program in Example 11.2 under the same conditions as [Using multithreading for concurrent activity](#) gives the following result:

Example 11.3 Result of *time python worker.py*

```

1 $ time python workerm.py
2 Results: [50495989, 50566997, 50474532, 50531418, 50522470, 50488087,  ←
   50498016, 50537899]
3 python workerm.py 16.53s user 0.12s system 363% cpu 4.581 total

```

The execution time has been reduced by 60 percent. Moreover, this time we've been able to consume up to 363 percent of CPU power, which is more than 90 percent (363/400) of the computer's CPU capacity.

Each time you think that you can parallelize some work, it's almost always better to rely on multiprocessing and to fork your jobs in order to spread the workload across several CPU cores. This wouldn't be a good solution for very small execution times, as the cost of the `fork()` call would be too big, but for larger computing needs, it works well.

11.3 Event-Driven Architecture

Event-driven programming is characterized by the use of events, such as user input, to dictate how control flows through a program, and it is a good solution for organizing

program flow. The event-driven program listens for various events happening on a queue and reacts based on those incoming events.

Let's say you want to build an application that listens for a connection on a socket and then processes the connection it receives. There are basically three ways to approach the problem:

1. Fork a new process each time a new connection is established, relying on something like the multiprocessing module.
2. Start a new thread each time a new connection is established, relying on something like the threading module.
3. Add this new connection to your event loop and react to the event it will generate when it occurs.

The problem of how a modern computer should handle tens of thousands of connections at the same time is known as **the C10K problem**. Among other things, the C10K resolution strategies explain how using an event loop to listen to hundreds of event sources is going to scale much better than, say, a one-thread-per-connection approach. This doesn't mean that the two techniques are not compatible, but it does mean that you can usually replace the multiple-threads approach with an event-driven mechanism.

Event-driven architecture uses an event loop: the program calls a function that blocks execution until an event is received and ready to be processed. The idea is that your program can be kept busy doing other tasks while waiting for inputs and outputs to complete. The most basic events are "Data ready to be read" and "Data ready to be written."

In Unix, the standard functions for building such an event loop are the system calls `select(2)` or `poll(2)`. These functions expect a list of file descriptors to listen for, and they will return as soon as at least one of the file descriptors is ready to be read from or written to.

In Python, we can access these system calls through the `select` module. It's easy enough to build an event-driven system with these calls, though doing so can be tedious. Example 11.4 shows an event-driven system that does our specified task: listening on a socket and processing any connections it receives.

Example 11.4 Event-driven program that listens for and processes connections

```

1  import select
2  import socket
3
4  server = socket.socket(socket.AF_INET,
5                          socket.SOCK_STREAM)
6  # Never block on read/write operations
7  server.setblocking(0)
8
9  # Bind the socket to the port
10 server.bind(('localhost', 10000))
11 server.listen(8)
12
13 while True:
14     # select() returns 3 arrays containing the object (sockets, files ←
15     ... ) that
16     # are ready to be read, written to or raised an error
17     inputs, outputs, excepts = select.select([server], [], [server])
18     if server in inputs:
19         connection, client_address = server.accept()
20         connection.send("hello!\n")

```

In Example 11.4, a server socket is created and set to non-blocking. That means that any read or write operation attempted on that socket will not block the program. If the program tries to read from the socket when there is no data ready to be read, the socket `recv()` method will raise an `OSError` indicating that the socket is not ready. If we did not call `setblocking(0)`, the socket would stay in blocking mode rather than raise an error, which is not what we want here. The socket is then bound to a port and listens with a maximum backlog of eight connections.

The main loop is built using `select()`, which receives the list of file descriptors we want to read (the socket in this case), the list of file descriptors we want to write to (none in this case), and the list of file descriptors we want to get exceptions from (the socket in this case). The `select()` function returns as soon as one of the selected file descriptors is ready to read, is ready to write, or has raised an exception. The returned values are lists of file descriptors positively matching the requests. It's then easy to check whether our socket is in the ready-to-be-read list and, if so, accept the connection and send a message.

11.4 Other Options and `asyncio`

Alternatively, there are many frameworks, such as `Twisted` or `Tornado`. `Twisted` that provide this kind of functionality in a more integrated manner; `Twisted` has been a de facto standard for years in this regard. C libraries that export Python interfaces, such as `libevent`, `libev` or `libuv`, also provides very efficient event loops.

These options all solve the same problem. The downside is that, while there are a wide variety of choices, most of them are not interoperable. Many are also callback based, meaning that the program flow is not very clear when reading the code; you have to jump to a lot of different places to read through the program.

Another option would be the `gevent` or `greenlet` libraries, which avoid the use of callback. However, the implementation details include CPython x86-specific code and dynamic modification of standard functions at runtime, meaning you wouldn't want to use and maintain code using these libraries over the long term.

In 2012, Guido Van Rossum began work on a solution code-named *tulip*, documented under `PEP 3156`. The goal of this package was to provide a standard event loop interface that would be compatible with all frameworks and libraries and be interoperable.

The *tulip* code has since been renamed and merged into Python 3.4 as the

asyncio module, and it is now the de facto standard. Not all libraries are compatible with asyncio, and most existing bindings need to be rewritten.

As of Python 3.6, asyncio has been so well integrated that it has its own `await` and `async` keywords, making it straightforward to use. Example 11.5 shows how the `aiohttp` library, which provides an asynchronous HTTP binding, can be used with asyncio to run several web page retrievals concurrently.

Example 11.5 Retrieving web pages concurrently with `aiohttp`

```

1  import aiohttp
2  import asyncio
3
4
5  async def get(url):
6      async with aiohttp.ClientSession() as session:
7          async with session.get(url) as response:
8              return response
9
10
11 loop = asyncio.get_event_loop()
12
13 coroutines = [get("http://example.com") for _ in range(8)]
14
15 results = loop.run_until_complete(asyncio.gather(*coroutines))
16
17 print("Results: %s" % results)

```

We define the `get()` function as asynchronous, so it is technically a coroutine. The two steps that compose the `get()` function, the connection and the page retrieval, are defined as asynchronous operations that yield control to the caller until they are ready. That makes it possible for asyncio to schedule another coroutine at any point. The module resumes the execution of a coroutine when the connection is established or the page is ready to be read. The eight coroutines are started and provided to the event loop at the same time, and it is `asyncio`'s job to schedule them efficiently.

The `asyncio` module is a great framework for writing asynchronous code and leveraging event loops. It supports files, sockets, and more, and a lot of third-party libraries are available to support various protocols. Don't hesitate to use it!

11.5 Service-Oriented Architecture

Python's shortcoming with regard to scaling for usage in large, complex applications can seem tricky to circumvent. However, Python is very good at implementing *service-oriented architecture* (SOA), a style of software design in which a set of services is provided by different components through a communication protocol. For example, SOA is the architecture type used by OpenStack in all its components. The components use HTTP REST to communicate with external clients (end users) and an abstracted remote procedure call (RPC) mechanism that is built on top of the Advanced Message Queuing Protocol (AMQP).

In your development situations, knowing which communication channels to use between those blocks is mainly a matter of knowing with whom you will be communicating.

When exposing a service to the outside world, the preferred channel is HTTP, especially for stateless designs such as REST-style (REpresentational State Transfer-style) architectures. These kinds of architectures make it easier to implement, scale, deploy, and comprehend services.

However, when exposing and using your API internally, HTTP may be not the best protocol. There are many other communication protocols, and a full description of any of them would likely fill an entire book.

In Python, there's plenty of libraries to build RPC systems. **Kombu** – among others – is interesting because it provides an RPC mechanism on top of a lot of back-ends, **AMQ protocol** being the main one. Support for **Redis**, **MongoDB**, **BeanStalk**, **Amazon SQS**, **CouchDB**, or **ZooKeeper** is also provided.

In the end, there's a huge amount of performance to be gained indirectly from using such loosely coupled architecture. If we consider that each module provides and exposes an API, we can run multiple daemons that can also expose that API, allowing multiple processes — and therefore CPUs — to handle the workload. For example, *Apache httpd* would create a new worker using a new system process that handles new connections; we could then dispatch a connection to a different worker running on the same node. To do so, we just need a system for dispatching the work to our various workers, which this API provides. Each block will be a different Python process, and as we've seen previously, this approach is better than multithreading for spreading out your workload. You'll be able to start multiple workers on each node. Even if stateless blocks are not strictly necessary, you should favor their use anytime you have the choice.

11.6 Interprocess Communication with ZeroMQ

As we've just discussed, a messaging bus is always needed when building distributed systems. Your processes need to communicate with each other in order to pass messages. [ZeroMQ](#) is a socket library that can act as a concurrency framework. [Workers using ZeroMQ](#) implements the same worker seen in [Example 11.2](#) but uses ZeroMQ as a way to dispatch work and communicate between processes.

Workers using ZeroMQ

```

1 import multiprocessing
2 import random
3 import zmq
4
5
6 def compute():
7     return sum(
8         [random.randint(1, 100) for i in range(1000000)]
9     )

```

```

10
11 def worker():
12     context = zmq.Context()
13     work_receiver = context.socket(zmq.PULL)
14     work_receiver.connect("tcp://0.0.0.0:5555")
15     result_sender = context.socket(zmq.PUSH)
16     result_sender.connect("tcp://0.0.0.0:5556")
17     poller = zmq.Poller()
18     poller.register(work_receiver, zmq.POLLIN)
19
20     while True:
21         socks = dict(poller.poll())
22         if socks.get(work_receiver) == zmq.POLLIN:
23             obj = work_receiver.recv_pyobj()
24             result_sender.send_pyobj(obj())
25
26
27 context = zmq.Context()
28 # Build a channel to send work to be done
29 work_sender = context.socket(zmq.PUSH)
30 work_sender.bind("tcp://0.0.0.0:5555")
31 # Build a channel to receive computed results
32 result_receiver = context.socket(zmq.PULL)
33 result_receiver.bind("tcp://0.0.0.0:5556")
34 # Start 8 workers
35 processes = []
36 for x in range(8):
37     p = multiprocessing.Process(target=worker)
38     p.start()
39     processes.append(p)
40 # Send 8 jobs
41 for x in range(8):
42     work_sender.send_pyobj(compute)
43 # Read 8 results
44 results = []
45 for x in range(8):
46     results.append(result_receiver.recv_pyobj())

```

```

47 # Terminate all processes
48 for p in processes:
49     p.terminate()
50 print("Results: %s" % results)

```

We create two sockets, one to send the function (`work_sender`) and one to receive the job (`result_receiver`). Each worker started by `multiprocessing.Process` creates its own set of sockets and connects them to the master process. The worker then executes whatever function is sent to it and sends back the result. The master process just has to send eight jobs over its sender socket and wait for eight results to be sent back via the receiver socket□ .

As you can see, ZeroMQ provides an easy way to build communication channels. I've chosen to use the TCP transport layer here to illustrate the fact that we could run this over a network. It should be noted that ZeroMQ also provides an interprocess communication channel that works locally (without any network layer involved) by using Unix sockets. Obviously, the communication protocol built upon ZeroMQ in this example is very simple for the sake of being clear and concise, but it shouldn't be hard to imagine building a more sophisticated communication layer on top of it. It's also easy to imagine building an entirely distributed application communication with a network message bus such as ZeroMQ or AMQP.

Note that protocols such as HTTP, ZeroMQ, and AMQP are language agnostic: you can use different languages and platforms to implement each part of your system. While we all agree that Python is a good language, other teams might have other preferences, or another language might be a better solution for some part of a problem.

In the end, using a transport bus to decouple your application into several parts is a good option. This approach allows you to build both synchronous and asynchronous APIs that can be distributed from one computer to several thousand. It doesn't tie you to a particular technology or language, so you can evolve everything in the right direction.

11.7 Summary

The rule of thumb in Python is to use threads only for I/O-intensive workload and to switch to multiple processes as soon as a CPU-intensive workload is on the table. Distributing workload on a wider scale — such as when building a distributed system over a network — requires external libraries and protocols. These are supported by Python, though provided externally.

As I wrote at the beginning of this chapter, scalability, distributed systems, and service-oriented architecture deserve an entire book. My second book, [Scaling Python](#), is dedicated entirely to writing distributed and scalable Python applications.

CHAPTER 12



Managing Relational Databases

Applications will almost always have to store data of some kind, and developers will often use a relational database management system (RDBMS) for that storage, combined with some type of object relational mapping tool (ORM). ORMs and RDBMSs can be tricky and are not a favorite topic for many developers, but sooner or later they must be confronted.

12.1 RDBMS, ORMs, and When to Use Them

An RDBMS is the database that stores the relational data of an application. Developers will use a language like SQL (Structured Query Language) to deal with the relational algebra, meaning that a language like this handles the management of the data and the relationships between the data. Used together, they allow you to both store data and query that data to get specific information as efficiently as possible. Having a good understanding of the structure of relational databases, such as how to use proper normalization or the different types of serializability, might keep you from falling into many traps. Obviously, such subjects deserve an entire book and

won't be covered in their entirety in this chapter; instead we'll focus on using the database via its usual programming language, SQL.

Developers may not want to invest in learning a whole new programming language to interact with the RDBMS. If so, they tend to avoid writing SQL queries entirely, relying instead on a library to do the work for them. ORM libraries are commonly found in programming language ecosystems, and Python is no exception.

The purpose of an ORM is to make database systems easier to access by abstracting the process of creating queries: generating the SQL so you don't have to. Unfortunately, this abstraction layer can prevent you from performing more specific or low-level tasks that the ORM is simply not capable of doing, such as writing complex queries.

There is also a particular set of difficulties with using ORMs in object-oriented programs that are so common, they are known collectively as the **object-relational impedance mismatch**. This impedance mismatch occurs because relational databases and object-oriented programs have different representations of data that don't map properly to one another: mapping SQL tables to Python classes won't give you optimal results, no matter what you do.

Having an understanding of SQL and RDBMSs will allow you to write your own queries directly, without having to rely on the abstraction layer for everything.

But that's not to say you should avoid ORMs entirely. ORM libraries can help with rapid prototyping of your application model, and some even provide useful tools such as schema upgrades and downgrades. The important thing is that you understand that using an ORM is not a substitute for gaining a real understanding of RDBMSs: many developers try to solve problems in the language of their choice rather than using their model API, and the solutions they come up with are inelegant at best.

Note

This chapter assumes you know basic SQL. Introducing SQL queries and how tables work is beyond the scope of this book. If you're new to SQL, I recommend learning the basics before continuing. *Practical SQL* by Anthony DeBarros (No Starch Press, 2018) is a good place to start. If you're into PostgreSQL, *Mastering PostgreSQL* by Dimitri Fontaine is a must-have.

To understand why having a proper understanding of RDBMS can help you write better code, let me give you an example. Say you have an SQL table for keeping track of messages. This table has a single column named `id` representing the ID of the message sender, which is the primary key, and a string containing the content of the message, like so:

```
1 CREATE TABLE message (  
2   id serial PRIMARY KEY,  
3   content text  
4 );
```

We want to detect any duplicate messages received and exclude them from the database. To do this, a typical developer might write SQL like this using an ORM:

Example 12.1 Detecting and excluding duplicate messages with an ORM

```
1 if query.select(Message).filter(Message.id == some_id):  
2     # We already have the message, it's a duplicate, ignore and raise  
3     raise DuplicateMessage(message)  
4 else:  
5     # Insert the message  
6     query.insert(message)
```

This code would work for most cases, but it has some major drawbacks:

- The duplicate constraint is already expressed in the SQL schema, so there is a sort

12.1. RDBMS, ORMS, AND WHEN TO USE THEM

of code duplication: using `PRIMARY KEY` implicitly defines the uniqueness of the `id` field.

- If the message is not yet in the database, this code executes two `SQL` queries: a `SELECT` statement and then an `INSERT` statement. Executing an `SQL` query might be long and require a round-trip to the `SQL` server, introducing extraneous delay.
- It doesn't take into account the possibility that someone else might insert a duplicate message after we call `select_by_id()` but before we call `insert()`, which would cause the program to raise an exception. This vulnerability is called a race condition.

There's a much better way to write this code, but it requires cooperation with the `RDBMS` server. Rather than checking for the message's existence and then inserting it, we can insert it right away and use a `try/except` block to catch a duplicate conflict:

```
1 try:
2     # Insert the message
3     message_table.insert(message)
4 except UniqueViolationError:
5     # Duplicate
6     raise DuplicateMessage(message)
```

In this case, inserting the message directly into the table works flawlessly if the message is not already present. If it is, the `ORM` will raise an exception indicating the violation of the uniqueness constraint. This method achieves the same effect as Example 12.1 but in a more efficient fashion and without any race condition. This is a very simple pattern, and it doesn't conflict with any `ORM` in any way. The problem is that developers tend to treat `SQL` databases as dumb storage rather than as a tool they can use to get proper integrity and consistency of data; consequently, they may duplicate the constraints written in `SQL` in their controller code rather than in their model.

Treating your `SQL` backend as a model API is good way to make efficient use

of it. You can manipulate the data stored in your RDBMS with simple function calls programmed in its own procedural language.

12.2 Database Backends

Another point that needs to be raised about ORM is support for multiple database backends. No ORM library provides a complete abstraction of all RDBMS features, and simplifying the code to the most basic RDBMS available will using any advanced RDBMS functions impossible without breaking the abstraction layer. Even simple things that aren't standardized in SQL, such as handling timestamp operations, are a pain to deal with when using an ORM. This is even truer if your code is RDBMS agnostic. It is important to keep this in mind when you choose the RDBMS for your application.

A good way to mitigate the problems with ORM libraries is to isolate them as prescribed in Section 2.3. This approach not only allows you to easily swap your ORM library for a different one should the need arise but also allows you to optimize your SQL usage by identifying places with inefficient usage of queries, bypassing most of the ORM boilerplate.

For example, an easy way to build such isolation is to use your ORM in a module of your application, for example `myapp.storage`. This module should export only functions and methods that allow you to manipulate the data at a high level of abstraction. The ORM should be used only from that module. At any point, you will be able to drop in any module providing the same API to replace `myapp.storage`.

The most commonly used ORM library in Python (and arguably the de facto standard) is [SQLAlchemy](#). This library supports a huge number of backends and provides abstraction for most common operations. Schema upgrades can be handled by third-party packages such as [alembic](#).

Some frameworks, such as [Django](#), provide their own ORM libraries. If you

choose to use a framework, it's smart to use the built-in library, which will have better integration with the framework than will an external one.

Warning



The Module View Controller (MVC) architecture that most frameworks rely on can be easily misused. These frameworks implement (or make it easy to implement) ORM in their models directly, but without abstracting enough of it: any code you have in your view and controllers that uses the model will also be using ORM directly. You need to avoid this. You should write a data model that includes the ORM library rather than consists of it: doing so will provide better testability and better isolation, as well as make it easier to swap out the ORM with another storage technology should the need arise.

12.3 Streaming data with Flask and PostgreSQL

Here, I'll show you how you can use one of PostgreSQL's advanced features to build an HTTP event-streaming system to help master your data storage.

12.3.1 Writing the Data-Streaming Application

The purpose of the micro-application in Example 12.2 to store messages in an SQL table and provide access to those messages via an HTTP REST API. Each message consists of a channel number, a source string, and a content string.

Example 12.2 SQL table schema for storing messages

```

1 CREATE TABLE message (
2   id SERIAL PRIMARY KEY,
3   channel INTEGER NOT NULL,
4   source TEXT NOT NULL,
5   content TEXT NOT NULL
6 );

```

We also want to stream these messages to the client so that it can process them in real time. To do this, we're going to use the LISTEN and NOTIFY features of PostgreSQL. These features allow us to listen for messages sent by a function we provide that PostgreSQL will execute:

Example 12.3 The *notify_on_insert* function

```

1 CREATE OR REPLACE FUNCTION notify_on_insert() RETURNS trigger AS $$
2 BEGIN
3   PERFORM pg_notify('channel_' || NEW.channel,
4                     CAST(row_to_json(NEW) AS TEXT));
5   RETURN NULL;
6 END;
7 $$ LANGUAGE plpgsql;

```

This code creates a trigger function written in *pl/pgsql*, a language that only PostgreSQL understands. Note that we could also write this function in other languages, such as Python itself, as PostgreSQL embeds the Python interpreter in order to provide a *pl/python* language. The single and simple operation we'll be performing here does not necessitate leveraging Python, so sticking to *pl/pgsql* is a wise choice.

The function `notify_on_insert()` performs a call to `pg_notify()`, which is the function that actually sends the notification. The first argument is a string that represents a channel, while the second is a string carrying the actual payload. We define the channel dynamically based on the value of the channel column in the row. In this case, the payload will be the entire row in JSON format. Yes, PostgreSQL knows how to convert a row to JSON natively!

12.3. STREAMING DATA WITH FLASK AND POSTGRESQL

Next we want to send a notification message on each INSERT performed in the message table, so we need to trigger this function on such events:

Example 12.4 The trigger for *notify_on_insert*

```
1 CREATE TRIGGER notify_on_message_insert AFTER INSERT ON message
2 FOR EACH ROW EXECUTE PROCEDURE notify_on_insert();
```

The function is now plugged in and will be executed upon each successful INSERT performed in the message table.

We can check that it works by using the LISTEN operation in psql:

```
1 $ psql
2 psql (9.3rc1)
3 SSL connection (cipher: DHE-RSA-AES256-SHA, bits: 256)
4 Type "help" for help.
5
6 mydatabase=> LISTEN channel_1;
7 LISTEN
8 mydatabase=> INSERT INTO message(channel, source, content)
9 mydatabase-> VALUES(1, 'jd', 'hello world');
10 INSERT 0 1
11 Asynchronous notification "channel_1" with payload
12 "{\"id\":1,\"channel\":1,\"source\":\"jd\",\"content\":\"hello world\"}"
13 received from server process with PID 26393.
```

As soon as the row is inserted, the notification is sent, and we're able to receive it through the PostgreSQL client. Now all we have to do is build the Python application that streams this event, shown in Example 12.5.

Example 12.5 Receiving notifications in Python

```
1 import psycopg2
2 import psycopg2.extensions
3 import select
4
5 conn = psycopg2.connect(database='mydatabase', user='myuser',
```

12.3. STREAMING DATA WITH FLASK AND POSTGRESQL

```
6         password='idkfa', host='localhost')
7
8 conn.set_isolation_level(
9     psycopg2.extensions.ISOLATION_LEVEL_AUTOCOMMIT)
10
11 curs = conn.cursor()
12 curs.execute("LISTEN channel_1;")
13
14 while True:
15     select.select([conn], [], [])
16     conn.poll()
17     while conn.notifies:
18         notify = conn.notifies.pop()
19         print("Got NOTIFY:", notify.pid, notify.channel, notify.payload ←
20             )
```

Example 12.5 connects to PostgreSQL using the `psycopg2` library. The `psycopg2` library is a Python module that implements the PostgreSQL network protocol and allows us to connect to a PostgreSQL server to send SQL requests and receive results. We could have used a library that provides an abstraction layer, such as `SQLAlchemy`, but abstracted libraries don't provide access to the `LISTEN` and `NOTIFY` functionality of PostgreSQL. It's important to note that it is still possible to access the underlying database connection to execute the code when using a library like `SQLAlchemy`, but there would be no point in doing that for this example, since we don't need any of the other features the ORM library would provide.

The program listens on `channel_1`, and as soon as it receives a notification, prints it to the screen. If we run the program and insert a row in the message table, we get the following output:

```
1 $ python listen.py
2 Got NOTIFY: 28797 channel_1
3 {"id":10,"channel":1,"source":"jd","content":"hello world"}
```

As soon as we insert the row, PostgreSQL runs the trigger and sends a notification. Our program receives it and prints the notification payload; here, that's the

row serialized to JSON. We now have the basic ability to receive data as it is inserted into the database, without doing any extra requests or work.

12.3.2 Building the Application

Next, we'll use **Flask**, a simple HTTP micro-framework, to build our application. We're going to build an HTTP server that streams the flux of insert using the *Server-Sent Events* message protocol defined by HTML5. An alternative would be to use Transfer-Encoding:chunked defined by HTTP/1.1.

Example 12.6 Flask streamer application

```

1 import flask
2 import psycopg2
3 import psycopg2.extensions
4 import select
5
6
7 app = flask.Flask(__name__)
8
9
10 def stream_messages(channel):
11     conn = psycopg2.connect(database='mydatabase', user='mydatabase',
12                             password='mydatabase', host='localhost')
13     conn.set_isolation_level(
14         psycopg2.extensions.ISOLATION_LEVEL_AUTOCOMMIT)
15
16     curs = conn.cursor()
17     curs.execute("LISTEN channel_%d;" % int(channel))
18
19     while True:
20         select.select([conn], [], [])
21         conn.poll()
22         while conn.notifies:
23             notify = conn.notifies.pop()

```

```

24         yield "data: " + notify.payload + "\n\n"
25
26
27 @app.route("/message/<channel>", methods=['GET'])
28 def get_messages(channel):
29     return flask.Response(stream_messages(channel),
30                           mimetype='text/event-stream')
31
32
33 if __name__ == "__main__":
34     app.run()

```

This application is simple enough that it supports streaming but not any other data retrieval operation. We use Flask to route the HTTP request GET /message/channel to our streaming code. As soon as the code is called, the application returns a response with the mimetype text/event-stream and sends back a generator function instead of a string. Flask will call this function and send results each time the generator yields something.

The generator, `stream_messages()`, reuses the code we wrote earlier to listen to PostgreSQL notifications. It receives the channel identifier as an argument, listens to that channel, and then yields the payload. Remember that we used PostgreSQL's JSON encoding function in the trigger function, so we're already receiving JSON data from PostgreSQL. There's no need for us to transcode the data since it's fine to send JSON data to the HTTP client.



Note

For the sake of simplicity, this example application has been written in a single file. If this were a real application, I would move the storage-handling implementation into its own Python module.

We can now run the server:

12.3. STREAMING DATA WITH FLASK AND POSTGRESQL

```
1 $ python listen+http.py
2 * Running on http://127.0.0.1:5000/
```

On another terminal, we can connect and retrieve the events as they're entered. Upon connection, no data is received, and the connection is kept open:

```
1 $ curl -v http://127.0.0.1:5000/message/1
2 * About to connect() to 127.0.0.1 port 5000 (#0)
3 * Trying 127.0.0.1...
4 * Adding handle: conn: 0x1d46e90
5 * Adding handle: send: 0
6 * Adding handle: recv: 0
7 * Curl_addHandleToPipeline: length: 1
8 * - Conn 0 (0x1d46e90) send_pipe: 1, recv_pipe: 0
9 * Connected to 127.0.0.1 (127.0.0.1) port 5000 (#0)
10 > GET /message/1 HTTP/1.1
11 > User-Agent: curl/7.32.0
12 > Host: 127.0.0.1:5000
13 > Accept: */*
14 >
```

But as soon as we insert some rows in the message table, we'll start seeing data coming in through the terminal running curl. In a third terminal, we insert a message in the database:

```
1 mydatabase=> INSERT INTO message(channel, source, content)
2 mydatabase-> VALUES(1, 'jd', 'hello world');
3 INSERT 0 1
4 mydatabase=> INSERT INTO message(channel, source, content)
5 mydatabase-> VALUES(1, 'jd', 'it works');
6 INSERT 0 1
```

Here's the data output:

```
1 data: {"id":71,"channel":1,"source":"jd","content":"hello world"}
2
3 data: {"id":72,"channel":1,"source":"jd","content":"it works"}
```

This data is printed on the terminal running `curl`. This keeps `curl` connected to the HTTP server while it waits for the next flux of messages. We created a streaming service without doing any kind of polling here, building an entirely push-based system, where information flows from one point to another seamlessly.

A naive and arguably more portable implementation of this application would instead repeatedly loop over a `SELECT` statement to poll for new data inserted in the table. This would work with any other storage system that does not support a publish-subscribe pattern as this one does.

12.4 Dimitri Fontaine on Databases

Dimitri Fontaine is a skilled PostgreSQL Major Contributor who works at Citus Data and argues with other database gurus on the *pgsql-hackers* mailing list. We've shared a lot of open source adventures, and he's been kind enough to answer some questions about what you should do when dealing with databases.

“



What advice would you give to developers using RDBMS as their storage backends?

RDBMSs were invented in the '70s to solve some common problems plaguing every application developer at that time, and the main services implemented by RDBMS were not simply data storage.

The main services offered by a RDBMS are actually the following:

- **Concurrency:** Access your data for read or write with as many concurrent threads of execution as you want — the RDBMS is there to handle that correctly for you. That's the main feature you want out of an RDBMS.
- **Concurrency semantics:** The details about the concurrency behavior when using a RDBMS are proposed with a high-level specification in terms of atomicity and isolation, which are maybe the most crucial parts of *ACID* (atomicity, consistency, isolation, durability). Atomicity is the property that between the time you BEGIN a transaction and the time you're done with it (either COMMIT or ROLLBACK), no other concurrent activity on the system is allowed to know what you're doing, whatever that is. When using a proper RDBMS, also include the Data Definition Language (*DDL*), for example, CREATE TABLE or ALTER TABLE. Isolation is all about what you're allowed to notice of the concurrent activity of the system from within your own transaction. The SQL standard defines four levels of isolation, as described in the [PostgreSQL documentation](#).

The RDBMS takes full responsibility for your data. So it allows the developer to describe their own rules for consistency, and then it will check that those rules are valid at crucial times such as at transaction commit or at statement boundaries, depending on the deferability of your constraints declarations.

The first constraint you can place on your data is its expected input and output formatting, using the proper data type. An RDBMS will know how to work with much more than text, numbers, and dates and will properly handle dates that actually appear in a calendar in use today.

Data types are not just about input and output formats, though. They also allow to implement behaviors and some level of polymorphism, as we all expect the basic equality tests to be data type specific: we don't compare text and numbers, dates and IP addresses, arrays and ranges, and so on in the same way.

Protecting your data also means that the only choice for a RDBMS is to actively refuse data that won't match your consistency rules, the first of which is the data type you've chosen. If you think it's okay to have to deal with a date such as 0000-00-00 that never existed in the calendar, then you need to rethink.

The other part of the consistency guarantees is expressed in terms of constraints as in CHECK constraints, NOT NULL constraints and constraint triggers, one of which is known as foreign key. All of that can be thought of as a user-level extension of the data type definition and behavior, the main difference being that you can choose to DEFER the enforcement of checking those constraints from the end of each statement to the end of the current transaction.

The relational bits of an RDBMS are all about modeling your data and the guarantee that all tuples found in a relation share a common set of rules: structure and constraints. When enforcing that, we are enforcing the use of a proper explicit schema to handle our data.

Working on a proper schema for your data is known as normalization, and you can aim for a number of subtly different normal forms in your design. Sometimes though, you need more flexibility than given by the result of your normalization process. Common wisdom is to first normalize your data schema and only then see about how to denormalize it in order to get back the flexibility you think you need. Chances are that you realize you actually don't need any.

When you do need more flexibility, then using PostgreSQL, you can

pick from a number of denormalization options: composite types, records, arrays, H-Store, JSON, or XML for starters.

There's a very important drawback to denormalization though, which is that the query language we're going to talk about next is designed to handle rather normalized data. With PostgreSQL, of course, the query language has been extended to support as much denormalization as possible when using composite types, arrays or H-Store, and even JSON in recent releases.

The RDBMS knows a lot about your data and can help you implement a very fine-grain security model, should you need to do so. The access patterns are managed at the relation and column level, and PostgreSQL also implements SECURITY DEFINER stored procedures, allowing you to offer access to sensible data in a very controlled way, much the same as with using saved user ID (SUID) programs.

The RDBMS offers to access your data using an SQL, which became a de facto standard in the '80s and is now driven by a committee. In the case of PostgreSQL, lots of extensions are being added, with each and every major release allowing you to have access to a very rich DSL language. All the work of query planning and optimization is done for you by the RDBMS so that you can focus on a declarative query where you describe only the result you want from the data you have.

And that's also why you need to pay close attention to the NoSQL offerings here, as most of those trendy products are in fact not removing just the SQL from the offering but a whole lot of other foundations that you've been trained to expect.

What advice would you give to developers using RDBMS as their storage backends?

My advice is to remember the differences between a storage backend

and an RDBMS. Those are very different services, and if all you need is a storage backend, maybe consider not using a RDBMS.

Most often, though, what you really need is a full-blown RDBMS. In that case, the best option you have is PostgreSQL. Go read its [documentation](#); see the list of data types, operators, functions, features, and extensions it provides. Read some usage examples on blog posts.

Then consider PostgreSQL a tool you can leverage in your development and include it in your application architecture. Parts of the services you need to implement are best offered at the RDBMS layer, and PostgreSQL excels at being that trustworthy part of your whole implementation.

What's the best way to use or not use an ORM?

The ORM will best work for CRUD applications: create, read, update, and delete. The read part should be limited to a very simple SELECT statement targeting a single table, as retrieving more columns than necessary has a significant impact on query performances and resources used.

Any column you retrieve from the RDBMS and that you end up not using is pure waste of precious resources, a first scalability killer. Even when your ORM is able to fetch only the data you're asking for, you still then have to somehow manage the exact list of columns you want in each situation, without using a simple abstract method that will automatically compute the fields list for you.

The create, update, and delete queries are simple INSERT, UPDATE, and DELETE statements. Many RDBMSs offer optimizations that are not leveraged by ORMs, such as returning data after an INSERT.

Furthermore, in the general case, a relation is either a table or the result of any query. It's common practice when using an ORM to

build a relational mapping between defined tables and some model classes, or some other helper stubs.

If you consider the whole SQL semantics in their generalities, then the relational mapper should really be able to map any query against a class. You would then presumably have to build a new class for each query you want to run.

The idea when applied to our case is that you trust your ORM to do a better job than you at writing efficient SQL queries, even when you're not giving it enough information to work out the exact set of data you are interested in.

It's true that at times, SQL can get quite complex, though you're not going to get anywhere near simplicity by using an API-to-SQL generator that you can't control.

However, there are two cases where you can relax and use your ORM, provided that you're willing to accept the following compromise: at a later point, you may need to edit your ORM usage out of your code base.

- **Time to market:** When you're really in a hurry and want to gain market share as soon as possible, the only way to get there is to release a first version of your application and idea. If your team is more proficient at using an ORM than handcrafting SQL queries, then by all means just do that. You have to realize, though, that as soon as you're successful with your application, one of the first scalability problems you will have to solve is going to be related to your ORM's producing really bad queries. Also, your usage of the ORM will have painted you into a corner and bad code design decisions. But if you're there, you're successful enough to spend some refactoring money and remove any dependency toward the ORM, right?

- **CRUD application:** This is the real thing, where you are only editing a single tuple at a time and you don't really care about performance, like for the basic admin application interface.

What are the pros of using PostgreSQL over other databases when working with Python?

Here are my top reasons for choosing PostgreSQL as a developer:

- **Community support:** The PostgreSQL community is vast and welcoming to new users, and folks will typically take the time to provide the best possible answer. The [mailing lists](#) are still the best way to communicate with the community.
- **Data integrity and durability:** Any data you send to PostgreSQL is safe in its definition and your ability to fetch it again later.
- **Data types, functions, operators, arrays, and ranges:** PostgreSQL has a very rich set of data types that come with a host of operators and functions. It's even possible to denormalize using arrays or JSON data types and still be able to write advanced queries, including joins, against those.
- **The planner and optimizer:** It's worth taking the time to understand how complex and powerful these are.
- **Transactional DDL:** It's possible to ROLLBACK almost any command. Try it now: just open your psql shell against a database you have and type in `BEGIN; DROP TABLE foo; ROLLBACK;`, where you replace `foo` with the name of a table that exists in your local instance. Amazing, right?
- **PL/Python (and others such as C, SQL, Javascript, or Lua):** You can run your own Python code on the server, right where the data is, so you don't have to fetch it over the network just to process it and then send it back in a query to do the next level of JOIN.

- Specific indexing (GiST, GIN, SP-GiST, partial and functional): You can create Python functions to process your data from within PostgreSQL, then index the result of calling that function. When you issue a query with a **WHERE** clause calling that function, it's called only once with the data from the query; then it's matched directly with the contents of the index.
- Extensions: such extensions include **hstore**, a full blown key value store with flexible indexing, **ltree** for indexing nested tags, **pg_trgm** as a poor man's full text search solution, that supports indexing regular expression searches and unanchored **LIKE** queries, **ip4r** for quick searches of an IP address in a range, and a lot more.
- Foreign Data Wrappers: the **foreign data wrappers** are a whole class of extensions, implementing the SQL/MED standard (Management of External Data). The idea is to embed a connection driver right into the PostgreSQL server then expose it through the **CREATE SERVER** command. PostgreSQL provides an API to **foreign data wrapper** authors that allows them to implement read and write access to the remote data, and also where clauses push-down for efficient joining capabilities. You can even use the advanced SQL capabilities of PostgreSQL against data that you maintain with another piece of technology!
- LISTEN/NOTIFY: PostgreSQL implements an asynchronous server-to-client protocol called LISTEN/NOTIFY. The application may receive unsolicited messages from the server when something interesting happens, for example an **UPDATE** of some data. The **NOTIFY** command accepts a data payload so that you can e.g. notify your cache application the object id's to purge when the object just has been removed or updated. Of course, the notification only happens if the transaction actually did a successful **COMMIT**.

- COPY Streaming protocol: PostgreSQL implements a **streaming** protocol and uses it to implement its fully integrated replication solution. Now, that protocol is quite easy to use from an application and allows impressive performance boosts. As soon as you're working on more than a dozen row at a time, sometimes before, thing about using COPY against a **temporary table** then issuing a single statement joining to that temporary table: PostgreSQL knows how to join against other tables in all data modifying statements (**insert**, **update**, **delete**), and batch operation usually are way faster.

”

CHAPTER 13



Write Less, Code More

In this final chapter, I've compiled a few of the more advanced features of Python that I use to write better code. These are not limited to the standard Python library. We'll cover how to make your code compatible for both Python 2 and 3, how to create a Lisp-like method dispatcher, context managers, and how to create a boilerplate for classes with the `attr` module.

13.1 Using Six for Python 2 and 3 Support

As you likely know, Python 3 breaks compatibility with Python 2 and shifts things around. However, the basics of the language haven't changed between versions, which makes it possible to implement forward and backward compatibility, creating a bridge between Python 2 and Python 3.

Luckily for us, this module already exists! It's called *six* — because $2 \times 3 = 6$.

The first useful thing provided by `six` is the `six.PY3` variable, which is a Boolean that indicates whether you are running Python 3 or not. This is the pivot variable for any of your code base that has two versions: one for Python 2 and one for Python 3.

However, be careful not to abuse it; scattering your code base with `if six.PY3` is going to make it difficult to read and understand.

When we discussed generators Section 8.2 on, we saw that Python 3 has a great property whereby iterable objects are returned instead of lists in various built-in functions, such as `map` or `filter`. Python 3 therefore got rid of methods like `dict.iteritems()`, which was the iterable version of `dict.items()` in Python 2, in favor of making `dict.items()` return an iterator rather than a list. Obviously, this change in methods and their return types can break your Python 2 code.

The `six` module provides `six.iteritems()` for such cases, which can be used to replace Python 2-specific code like this:

```
1 for k, v in mydict.iteritems():
2     print(k, v)
```

Using `six`, you would replace the `mydict.iteritems()` code with Python 2 and 3 compliant code like so:

```
1 import six
2
3 for k, v in six.iteritems(mydict):
4     print(k, v)
```

And *voilà*, both Python 2 and Python 3 compliance achieved in a snap! The `six.iteritems()` function will use either `dict.iteritems()` or `dict.items()` to return a generator, depending on the version of Python you're using. The `six` module provides a lot of similar helper functions that can make it easy to support multiple Python versions.

Another example would be the `six` solution to the `raise` keyword, whose syntax is different between Python 2 and Python 3. In Python 2, `raise` will accept multiple arguments, but in Python 3, `raise` accepts an exception as its only argument and nothing else. Writing a `raise` statement with two or three arguments in Python 3 would result in a `SyntaxError`.

The `six` module provides a workaround here in the form of the `six.reraise()` function, which allows you to reraise an exception in whichever version of Python you use.

13.1.1 Strings and Unicode

One of the primary enhancements in Python 3 is the ability to handle advanced encodings, which solved the string and unicode issues of Python 2. In Python 2, the basic type for string is `str`, which can only handle basic ASCII strings. The type `unicode`, added later in Python 2.5, handles real strings of text.

In Python 3, the basic string type is still `str`, but it shares the properties of the Python 2 `unicode` class and can handle advanced encodings. The `bytes` type replaces the `str` type for handling basic character streams.

The `six` module again provides functions and constants, such as `six.u` and `six.string_types`, to handle the transition. The same compatibility is provided for integers, with `six.integer_types` that will handle the `long` type that has been removed from Python 3.

13.1.2 Handling Python Modules Moves

In the Python Standard Library, some modules have moved or have been renamed between Python 2 and 3. The `six` module provides a module called `six.moves` that handles a lot of these moves transparently.

For example, the `ConfigParser` module from Python 2 has been renamed to `configparser` in Python 3. Example [13.1](#) shows how code can be ported and made compatible with both major Python versions using `six.moves`:

Example 13.1 Using `six.moves` to use `ConfigParser()` with Python 2 and Python 3

```
1 from six.moves.configparser import ConfigParser
2
3 conf = ConfigParser()
```

It is also possible to add your own moves via `six.add_move` to handle code transitions that are not handled natively by `six`.

In the event that the `six` library doesn't cover all your use cases, it may be worth building a compatibility module encapsulating `six` itself, thereby ensuring that you will be able to enhance the module to fit future versions of Python or dispose of (part of) it when you want to stop supporting a particular version of the language. Also note that `six` is open source and that you can contribute to it rather than maintain your own hacks!

13.1.3 The modernize Module

Lastly, there is a tool named `modernize` that uses the `six` module to “modernize” your code by porting it to Python 3, rather than simply converting Python 2 syntax to Python 3 syntax. This provides support for both Python 2 and Python 3. The `modernize` tool helps to get your port off to a strong start by carrying out most of the grunt work for you, making this tool a better choice than the standard 2to3 tool.

13.2 Using Python Like Lisp to Make a Single Dispatcher

I like to say that Python is a good subset of the Lisp programming language, and as time passes, I find this to be more and more true. The [PEP 443](#) proves that point:

it describes a way to dispatch generic functions in a similar manner to that provided by the Common Lisp Object System (CLOS).

If you're familiar with Lisp, this won't be news to you. The Lisp object system, which is one of the basic components of Common Lisp, provides a simple, efficient way to define and handle method dispatching. I'll show you how generic methods work in Lisp first.

13.2.1 Creating Generic Methods in Lisp

To begin with, let's define a few very simple classes, without any parent classes or attributes, in Lisp:

```
1 (defclass snare-drum ()
2   ())
3
4 (defclass cymbal ()
5   ())
6
7 (defclass stick ()
8   ())
9
10 (defclass brushes ()
11   ())
```

This defines the classes `snare-drum`, `cymbal`, `stick`, and `brushes` without any parent class or attributes. These classes compose a drum kit, and we can combine them to play sound. For this, we define a `play()` method that takes two arguments and returns a sound as a string:

```
1 (defgeneric play (instrument accessory)
2   (:documentation "Play sound with instrument and accessory."))
```

This only defines a generic method that isn't attached to any class and so cannot yet be called. At this stage, we've only informed the object system that the method is

13.2. USING PYTHON LIKE LISP TO MAKE A SINGLE DISPATCHER

generic and might be called with two arguments named `instrument` and `accessory`. In Example 13.2, we'll implement versions of this method that simulate playing our snare drum.

Example 13.2 Defining generic methods in Lisp, independent of classes

```
1 (defmethod play ((instrument snare-drum) (accessory stick))
2   "POC!")
3
4 (defmethod play ((instrument snare-drum) (accessory brushes))
5   "SHHHH!")
6
7 (defmethod play ((instrument symbol) (accessory brushes))
8   "FRCCCHHT!")
```

Now we've defined concrete methods in code. Each method takes two arguments: `instrument`, which is an instance of `snare-drum` or `cymbal`, and `accessory`, which is an instance of `stick` or `brushes`.

At this stage, you should see the first major difference between this system and the Python (or similar) object systems: the method isn't tied to any particular class. The methods are *generic*, and they can be implemented for any class.

Let's try it: we can call our `play()` method with some objects:

```
1 * (play (make-instance 'snare-drum) (make-instance 'stick))
2 "POC!"
3
4 * (play (make-instance 'snare-drum) (make-instance 'brushes))
5 "SHHHH!"
```

As you can see, which function is called depends on the class of the arguments — the object system *dispatches* the function calls to the right function for us, based on the type of the arguments we pass. If we call `play()` with an object whose classes do not have a method defined, an error will be thrown.

In Example 13.3, the `play()` method is called with a `cymbal` and a `stick` in-

13.2. USING PYTHON LIKE LISP TO MAKE A SINGLE DISPATCHER

stance; however, the `play()` method has never been defined for those arguments, so it raises an error:

Example 13.3 Calling a method with an unavailable signature

```
1 * (play (make-instance 'cymbal) (make-instance 'stick))
2 debugger invoked on a SIMPLE-ERROR in thread
3 #<THREAD "main thread" RUNNING {1002ADAF23}>:
4   There is no applicable method for the generic function
5     #<STANDARD-GENERIC-FUNCTION PLAY (2)>
6     when called with arguments
7     (#<CYMBAL {1002B801D3}> #<STICK {1002B82763}>).
8
9 Type HELP for debugger help, or (SB-EXT:EXIT) to exit from SBCL.
10
11 restarts (invokable by number or by possibly-abbreviated name):
12   0: [RETRY] Retry calling the generic function.
13   1: [ABORT] Exit debugger, returning to top level.
14
15 ((:METHOD NO-APPLICABLE-METHOD (T)) #<STANDARD-GENERIC-FUNCTION PLAY (2)>
    (#<CYMBAL {1002B801D3}> #<STICK {1002B82763}>)) [fast-method]
```

CLOS provides even more features, such as method inheritance or object-based dispatching, rather than using classes. If you're really curious about the many features CLOS provides, I suggest that you read [“A Brief Guide to CLOS”](#) by Jeff Dalton as a starter.

13.2.2 Generic Methods in Python

Python implements a simpler version of this workflow with the `singledispatch()` function, which has been distributed as part of the `functools` module since Python 3.4. In versions 2.6 to 3.3, the `singledispatch()` function is provided through the Python Package Index; for those eager to try it out, just run `pip install singledispatch`.

13.2. USING PYTHON LIKE LISP TO MAKE A SINGLE DISPATCHER

Example 13.4 shows a rough equivalent of the Lisp program we built in Example 13.2.

Example 13.4 A Python implementation of our drum kit

```
1 import functools
2
3 class SnareDrum(object): pass
4 class Cymbal(object): pass
5 class Stick(object): pass
6 class Brushes(object): pass
7
8 @functools.singledispatch
9 def play(instrument, accessory):
10     raise NotImplementedError("Cannot play these")
11
12 @play.register(SnareDrum)
13 def _(instrument, accessory):
14     if isinstance(accessory, Stick):
15         return "POC!"
16     if isinstance(accessory, Brushes):
17         return "SHHHH!"
18     raise NotImplementedError("Cannot play these")
19
20 @play.register(Cymbal)
21 def _(instrument, accessory):
22     if isinstance(accessory, Brushes):
23         return "FRCCCHHT!"
24     raise NotImplementedError("Cannot play these")
```

Example 13.4 defines our four classes and a base `play()` function that raises `NotImplementedError`, indicating that by default we don't know what to do.

We then write a specialized version of the `play()` function for a specific instrument, the `SnareDrum`. This function checks which accessory type has been passed and returns the appropriate sound or raises `NotImplementedError` again if the accessory isn't recognized.

13.2. USING PYTHON LIKE LISP TO MAKE A SINGLE DISPATCHER

If we run the program, it works as follows:

```
1 >>> play(SnareDrum(), Stick())
2 'POC!'
3 >>> play(SnareDrum(), Brushes())
4 'SHHHH!'
5 >>> play(Cymbal(), Stick())
6 Traceback (most recent call last):
7 NotImplementedError: Cannot play these
8 >>> play(SnareDrum(), Cymbal())
9 NotImplementedError: Cannot play these
```

The `singledispatch` module checks the class of the first argument passed and calls the appropriate version of the `play` function. For the `object` class, the first defined version of the function is always the one that is run. Therefore, if our instrument is an instance of a class that we did not register, this base function will be called.

As we saw in the Lisp version of the code, CLOS provides a multiple dispatcher that can dispatch based on the type of *any of the arguments* defined in the method prototype, not just the first one. The Python dispatcher is named `singledispatch` for a good reason: it only knows how to dispatch based on the first argument.

In addition, `singledispatch` offers no way to call the parent function directly. There is no equivalent of the Python `super()` function; you'll have to use various tricks to bypass this limitation.

While Python is improving its object system and dispatch mechanism, it still lacks a lot of the more advanced features that something like CLOS provides out of the box. That makes encountering `singledispatch` in the wild pretty rare. It's still interesting to know it exists, as you may end up implementing such a mechanism yourself at some point.

13.3 Context Managers

The `with` statement introduced in Python 2.6 is likely to remind old-time Lispers of the various `with-*` macros that are often used in that language. Python provides a similar-looking mechanism with the use of objects that implement the context management protocol.

If you never used the context management protocol, here's how it works: the code block contained inside the `with` statement is surrounded by two function calls. Those two calls are determined by the object being used in the `with` statement. Those objects are said to implement the context management protocol.

Objects like those returned by `open()` support this protocol; that's why you can write code along these lines:

```
1 with open("myfile", "r") as f:
2     line = f.readline()
```

The object returned by `open()` has two methods: one called `__enter__` and one called `__exit__`. These methods are called at the start of the `with` block and at the end of it, respectively.

A simple implementation of a context object is shown in Example 13.5.

Example 13.5 Simple implementation of a context object

```
1 class MyContext(object):
2     def __enter__(self):
3         pass
4
5     def __exit__(self, exc_type, exc_value, traceback):
6         pass
```

This implementation does not do anything, but it is valid and shows the signature of the methods that need to be defined to provide a class following the context protocol.

13.3. CONTEXT MANAGERS

The context management protocol might be appropriate to use when you identify the following pattern in your code, where it is expected that a call to method B must always be done after a call to A:

1. Call method A.
2. Execute some code.
3. Call method B.

The `open()` function illustrates this pattern well: the constructor that opens the file and allocates a file descriptor internally is method A. The `close()` method that releases the file descriptor corresponds to method B. Obviously, the `close()` function is always meant to be called after you instantiate the file object.

It can be tedious to implement this protocol manually, so the `contextlib` standard library provides the `contextmanager` decorator to make implementation easier. The `contextmanager` decorator should be used on a generator function. The `__enter__` and `__exit__` methods will be dynamically implemented for you based on the code that wraps the `yield` statement of the generator.

In Example 13.5, `MyContext` is defined as a context manager:

Example 13.6 Simplest usage of `contextlib.contextmanager`

```
1 import contextlib
2
3 @contextlib.contextmanager
4 def MyContext():
5     print("do something first")
6     yield
7     print("do something else")
8
9
10 with MyContext():
11     print("hello world")
```

13.3. CONTEXT MANAGERS

The code before the `yield` statement will be executed before the `with` statement body is run; the code after the `yield` statement will be executed once the body of the `with` statement is over. When run, this program outputs the following:

```
1 do something first
2 hello world
3 do something else
```

There are a couple of things to handle here though. First, it's possible to `yield` something inside our generator that can be used as part of the `with` block.

Example 13.7 shows how to `yield` a value to the caller. The keyword `as` is used to store this value in a variable.

Example 13.7 Defining a context manager yielding a value

```
1 import contextlib
2
3 @contextlib.contextmanager
4 def MyContext():
5     print("do something first")
6     yield 42
7     print("do something else")
8
9
10 with MyContext() as value:
11     print(value)
```

Example 13.7 shows how to `yield` a value to the caller. The keyword `as` is used to store this value in a variable. When executed, the code outputs the following:

```
1 do something first
2 42
3 do something else
```

When a context manager is used, you might need to handle exceptions that can be raised within the `with` code block. This can be done by surrounding the `yield` statement with a `try/except` block, as shown in Example 13.8.

Example 13.8 Handling exceptions in a context manager

```
1 import contextlib
2
3 @contextlib.contextmanager
4 def MyContext():
5     print("do something first")
6     try:
7         yield 42
8     finally:
9         print("do something else")
10
11
12 with MyContext() as value:
13     print("about to raise")
14     raise ValueError("let's try it")
15     print(value)
```

Here, a `ValueError` is raised at the beginning of the `with` code block; Python will propagate this error back to the context manager, and the `yield` statement will appear to raise the exception itself. We enclose the `yield` statement in `try` and `finally` to make sure the final `print()` is run.

When executed, Example 13.8 outputs the following:

```
1 do something first
2 about to raise
3 do something else
4 Traceback (most recent call last):
5   File "<stdin>", line 3, in <module>
6   ValueError: let's try it
```

As you can see, the error is raised back to the context manager, and the program is able to resume and finish execution because it ignored the exception using a `try/finally` block.

In some contexts, it can be useful to use several context managers at the same time, for example, when opening two files at the same time to copy their content, as shown in Example 13.9.

Example 13.9 Opening two files at the same time to copy content

```
1 with open("file1", "r") as source:
2     with open("file2", "w") as destination:
3         destination.write(source.read())
```

That being said, since the `with` statement supports multiple arguments, it's actually more efficient to write a version using a single `with`, as shown in Example 13.10.

Example 13.10 Opening two files at the same time with one `with` statement

```
1 with open("file1", "r") as source, open("file2", "w") as destination:
2     destination.write(source.read())
```

Context managers are extremely powerful design patterns that help to make sure your code flow is always correct, no matter what exception might occur. They can help to provide a consistent and clean programming interface in many situations in which code should be wrapped by other code and `contextlib.contextmanager`.

13.4 Less Boilerplate with `attr`

Writing classes in Python can be cumbersome. You'll often find yourself repeating just a few patterns because there are no other options. One of the most common examples, as illustrated in Example 13.11, is when initializing an object with a few attributes passed to the constructor:

Example 13.11 Common class initialization boilerplate

```

1 class Car(object):
2     def __init__(self, color, speed=0):
3         self.color = color
4         self.speed = speed

```

The process is always the same: you copy the value of the argument passed to the `__init__` function to a few attributes stored in the object. Sometimes you'll also have to check the value that is passed, compute a default, and so on.

Obviously, you also want your object to be represented correctly if printed, so you'll have to implement a `__repr__` method. There's a chance some of your classes are simple enough to be converted to dictionaries for serialization. Things become even more complicated when talking about comparison and hashability (the ability to use hash on an object and store it in a set).

In reality, most Python programmers do none of this, because the burden of writing all those checks and methods is too heavy, especially when you're not always sure you'll need them. For example, you might find that `__repr__` is useful in your program only that one time you're trying to debug or trace it and decide to print objects in the standard output — and nothing else.

The `attr` library aims to solve this issue in a straightforward manner by providing a generic boilerplate for all your classes and generating much of the code for you. You can install `attr` using `pip` with the command `pip install attr`. Get ready to enjoy!

Once installed, the `attr.s` decorator is your entry point into the wonderful world of `attr`. Use it above a class declaration and then use the function `attr.ib()` to declare attributes in your classes. Example 13.12 shows a way to rewrite Example 13.11 using `attr`.

Example 13.12 Declaring a class with `attr`

```

1 import attr

```

13.4. LESS BOILERPLATE WITH ATTR

```
2
3 @attr.s
4 class Car(object):
5     color = attr.ib()
6     speed = attr.ib(default=0)
```

When declared this way, the class automatically gains a few useful methods for free, such as `__repr__`, which is called to represent objects when they are printed on `stdout` in the Python interpreter:

```
1 >>> Car("blue")
2 Car(color='blue', speed=0)
```

This is a cleaner output than the default `__repr__` would have printed:

```
1 '<__main__.Car object at 0x104ba4cf8>'
```

You can also add more validation on your attributes by using the `validator` and `converter` keyword arguments.

Example 13.13 shows how the `attr.ib()` function can be used to declare an attribute with some constraints:

Example 13.13 Using `attr.ib()` with its `converter` argument

```
1 import attr
2
3 @attr.s
4 class Car(object):
5     color = attr.ib(converter=str)
6     speed = attr.ib(default=0)
7
8     @speed.validator
9     def speed_validator(self, attribute, value):
10         if value < 0:
11             raise ValueError("Value cannot be negative")
```

13.4. LESS BOILERPLATE WITH ATTR

The `converter` argument manages the conversion of whatever is passed to the constructor. The `validator()` function can be passed as an argument to `attr.ib()` or used as a decorator, as shown in Example 13.13.

The `attr` module provides a few validators of its own (for example, `attr.validators.instance_of()` to check the type of the attribute), so be sure to check them out before wasting your time building your own.

The `attr` module also provides tweaks to make your object hashable so it can be used in a set or a key of a dictionary: just pass `frozen=True` to `attr.s` to make the class instances immutable.

Example 13.14 shows how using the `frozen` parameter changes the behavior of the class:

Example 13.14 Using `frozen=True` with `attr`

```
1 >>> import attr
2 >>> @attr.s(frozen=True)
3 ... class Car(object):
4 ...     color = attr.ib()
5 ...
6 >>> {Car("blue"), Car("blue"), Car("red")}
7 {Car(color='red'), Car(color='blue')}
8 >>> Car("blue").color = "red"
9 attr.exceptions.FrozenInstanceError
```

Example 13.14 shows how using the `frozen` parameter changes the behavior of the `Car` class: it can be hashed and therefore stored in a set, but objects cannot be modified anymore.

In summary, `attr` provides the implementation for a ton of useful methods, thereby saving you from writing them yourself. I highly recommend leveraging `attr` for its efficiency when building your classes and modeling your software.

Index

—

`__getattribute__`, [186](#)

`__main__`, [106](#)

`__slots__`, [253](#)

A

`abc`, [56](#), [180](#)

Abstract Methods, [179](#), [181](#), [184](#)

Abstract Syntax Tree, [213](#)

`aiohttp`, [283](#)

`alembic`, [46](#), [293](#)

`all()`, [203](#)

AMQP, [284](#)

`any()`, [203](#)

API Versions, [75](#)

Architecture, [274](#)

 Event-Driven Architecture, [279](#)

 Service-Oriented Architecture, [284](#)

`array`, [270](#)

`as`, [38](#)

`ast`, [214](#)

`ast.dump()`, [214](#)

`ast.parse()`, [214](#)

`ast.walk()`, [219](#)

`asyncio`, [282](#), [283](#)

`attr`, [324](#)

`autodoc`, [68](#)

automation, [32](#)

`autosummary`, [69](#)

B

`bisect`, [57](#), [248](#)

`bisect.bisect()`, [249](#)

Buffer Protocol, [263](#)

C

C, [31](#)

C10K, [280](#)

Christophe de Vienne, [82](#)

Class Decorators, [168](#)

Class Methods, [178](#), [181](#)

`classmethod`, [179](#)

Clojure, [228](#)

Closure, [247](#)

Coding Style, [28](#)

`collections`, [57](#), [239](#)

`collections.Counter`, [239](#)

`collections.defaultdict()`, [239](#)

Common Lisp Object System, [313](#)

`compile()`, [215](#)

`conf.py`, [66](#)

Console Scripts, [115](#)

Context Managers, [319](#)

contextlib, [320](#)

Coverage, [145](#)

cProfile, [240](#), [272](#)

CPython, [232](#), [261](#), [273](#), [282](#)

D

Data Structure, [237](#)

Database, [289](#), [293](#)

datetime, [89](#), [92](#), [93](#), [95](#), [96](#)

datetime.datetime, [88](#), [94](#)

datetime.datetime.now(), [89](#)

datetime.datetime.utcnow(), [89](#)

datetime.tzinfo, [92](#)

datetutil.tz.gettz, [92](#)

datetutil.zoneinfo, [91](#)

dateutil, [90–92](#), [95](#), [96](#)

debtcollector, [80–82](#)

Decorators, [163–165](#), [167](#), [170](#)

 Class Decorators, [168](#)

dependencies, [48](#)

DeprecationWarning, [78](#), [81](#)

dict.get(), [237](#)

Dimitri Fontaine, [301](#)

dis, [243](#)

dis.dis, [243](#)

Disassembling, [243](#)

distlib, [101](#)

distribute, [100](#), [101](#)

distutils, [98](#), [99](#), [101](#)

distutils2, [100–102](#)

Django, [53](#)

doctest, [70](#)

Documentation, [63](#), [64](#)

Doug Hellmann, [53](#)

E

easy_install, [99](#)

Entry Points, [100](#), [112](#), [113](#), [117](#), [118](#)

entry_point_inspector, [113](#)

enumerate(), [202](#)

Event-Driven Architecture, [279](#)

External Library, [46](#)

F

filter(), [201](#)

first, [207](#)

Fixtures, [135](#), [138](#)

flake8, [31](#), [156](#), [221](#)

flake8-import-order, [31](#)

Flask, [53](#), [298](#)

Flycheck, [32](#)

Frameworks, [52](#), [61](#)

Functional Programming, [190](#), [200](#)

Functions

 Pure Functions, [191](#)

functools, [207](#), [260](#)

functools.lru_cache, [260](#)

functools.partial(), [208](#)

functools.update_wrapper(), [171](#)

G

Generators, [192](#), [193](#), [197](#)

gevent, [282](#)

Git, [28](#), [51](#)

GitHub, [157](#)

Global Interpreter Lock, [273](#), [275](#)

Gnocchi, [74](#), [138](#)

GNU Emacs, [32](#)

greenlet, [282](#)

H

Hooks, [120](#)

Hy, [40](#), [228](#), [262](#)

I

import, [36–38](#)

init.py, [25](#)

inspect, [173](#), [197](#)

inspect.isgeneratorfunction, [197](#)

iso8601, [95](#)

isoformat, [93](#)

ISO 8601, [97](#)

Iterator Protocol, [192](#)

itertools, [205](#), [210](#)

J

JIT, [34](#), [263](#)

Joshua Harlow, [32](#)

Jython, [232](#), [276](#)

K

KCacheGrind, [241](#)

L

lambda, [207](#)

Layout, [23](#)

Libraries, [36](#)

 External Library, [46](#)

 Standard Library, [43](#)

Lisp, [228](#), [313](#)

List comprehension, [199](#)

M

map(), [201](#)

Memoization, [259](#)

memory_profiler, [263](#)

memoryview, [265](#)

Mercurial, [28](#)

Method Resolution Order, [186](#)

Methods, [174](#)

 Abstract Methods, [179](#), [181](#), [184](#)

 Class Methods, [178](#), [181](#)

 Static Methods, [176](#), [181](#)

mock, [139](#)

mock.ANY, [142](#)

mock.patch(), [142](#)

Mocking, [139](#)

modernize, [313](#)

mro, [186](#)

Multiprocessing, [277](#)

multiprocessing, [278](#)

multiprocessing.Process, [287](#)

Multithreading, [274](#), [277](#)

N

namedtuple, [57](#), [252](#), [256](#)

Nick Coghlan, [123](#)

NodeTransformer, [219](#)

noqa, [31](#)

O

Object-Relational Mapper, [289](#), [290](#)

OpenStack, [17](#), [32](#), [46](#), [57](#), [103](#), [262](#)

operator, [211](#)

Optimizations, [236](#)

ORM, [305](#)

P

packaging, [100](#), [101](#)

Parallel Testing, [134](#)

Paul Tagliamonte, [231](#)

pbr, [103](#), [104](#), [118](#)

PendingDeprecationWarning, [78](#)

People

Christophe de Vienne, [82](#)

Dimitri Fontaine, [301](#)

Doug Hellmann, [53](#)

Joshua Harlow, [32](#)

Nick Coghlan, [123](#)

Paul Tagliamonte, [231](#)

Robert Collins, [159](#)

Victor Stinner, [270](#)

pep8, [29](#)

PEP 20, [272](#)

PEP 255, [192](#)

PEP 257, [85](#)

PEP 302, [40](#)

PEP 3118, [265](#)

PEP 3156, [282](#)

PEP 376, [104](#)

PEP 426, [123](#), [124](#)

PEP 427, [104](#)

PEP 440, [26](#)

PEP 443, [313](#)

PEP 454, [270](#)

PEP 495, [96](#)

PEP 513, [107](#)

PEP 7, [31](#)

PEP 8, [28–30](#)

Performances, [236](#)

pip, [49](#), [51](#), [98](#), [104](#)

pkg_resources, [115](#), [118](#)

PostgreSQL, [294](#), [307](#)

Pure Functions, [191](#)

pyflakes, [31](#)

pylint, [31](#)

PyPI, [49](#), [108](#), [109](#)

pyprof2calltree, [241](#)

PyPy, [232](#), [261](#), [263](#)

pytest, [127](#), [153](#)

pytest-xdist, [134](#)

pytest.fixture, [135](#), [138](#)

pytest.mark, [133](#)

pytest.skip, [130](#)

Python Versions, [21](#)

PYTHONPATH, [39](#)

Python 2, [22](#)

Python 3, [22](#), [46](#), [53](#), [59](#)

R

range, [194](#)

README.rst, [24](#), [65](#), [103](#)

reST, [64](#)

Robert Collins, [159](#)

S

Scalability, [274](#)

Scaling Python, [274](#), [333](#)

Scenarios, [138](#)

Semantic Versioning, [27](#)

Service-Oriented Architecture, [284](#)

set, [238](#)

setup.cfg, [24](#), [100](#), [102](#), [103](#)

setup.py, [24](#), [98](#), [99](#), [102](#), [103](#)

setuptools, [23](#), [99–101](#), [104](#), [107](#), [116](#)

Single Responsibility Principle, [59](#)

six, [310](#)

six.moves, [312](#)

Skipping Tests, [130](#)

Slots, [252](#)

smiley, [54](#)

sorted(), [202](#)

Sphinx, [63](#), [64](#), [66](#), [81](#)

autodoc, [68](#)

autosummary, [69](#)

doctest, [70](#)

Sphinx Extension, [72](#)

Sphinx Extension, [72](#)

sphinx-build, [66](#)

SQL, [289](#), [306](#)

SQLAlchemy, [46](#), [60](#), [293](#)

sqlalchemy-migrate, [46](#)

Standard Library, [43](#)

Static Methods, [176](#), [181](#)

staticmethod, [177](#)

stevedore, [55](#), [120](#)

StopIteration, [194](#)

str, [312](#)

struct, [270](#)

super, [185](#), [188](#)

sys, [38](#)

sys.meta_path, [40](#)

sys.modules, [38](#)

sys.path, [39](#)

sys.path_hooks, [40](#)

T

Taskflow, [32](#)

Testing

Coverage, [145](#)

Fixtures, [135](#), [138](#)

Mocking, [139](#)

Parallel Testing, [134](#)

Scenarios, [138](#)

Testing Policy, [157](#)

Unit Testing, [126](#)

Testing Policy, [157](#)

Tests

Skipping Tests, [130](#)

Time zones, [87](#), [88](#), [92](#)
timeit, [245](#), [272](#)
timestamp, [88](#)
tox, [152](#), [155](#), [156](#)
tox-travis, [158](#)
Travis CI, [157](#)
Twisted, [53](#)
tzinfo, [89–93](#)

U

unicode, [312](#)
Unit Testing, [126](#)
uprefix, [43](#)
UTC, [88](#)

V

Versions, [26](#)
 API Versions, [75](#)
 Python Versions, [21](#)
Victor Stinner, [270](#)
vim, [32](#)
Virtual Environments, [149](#)
virtualenv, [150](#)

W

Warnings, [80](#)
warnings, [78](#)
wheel, [104](#), [107](#)
with, [319](#)
WSGI, [58](#)
WSME, [86](#)

Y

yield, [193](#), [195](#)

Z

Zake, [32](#)
Zero Copy, [263](#)
ZeroMQ, [285](#)
zip(), [204](#)

Beyond This Book

Congratulations! You made it to the end of the book. You've just upped your Python game and increased your knowledge of writing in Python more efficiently and productively. I hope you enjoyed reading it as much as I enjoyed writing it.

Python is a wonderful language and can be used in many different fields, and there are many more areas of Python that we did not touch on in this book. But every book needs an ending, right?

If your hunger for knowledge is not yet satisfied, you can continue to learn about Python by digging through a few other interesting resources, specifically [Scaling Python](#), another book written by yours truly about writing network-distributed Python application at very large scale.

I also highly recommend profiting from open source by reading the available source code out there and contributing to it. Having your code reviewed and discussed by other developers is often a great way to learn.

Happy hacking!