



GO

FUNDAMENTALS

MARK BATES | CORY LANOU

GOPHER GUIDES



Foreword by STEVE FRANCIA

FREE SAMPLE CHAPTER |



Go Fundamentals

This page intentionally left blank

Go Fundamentals

Mark Bates
Cory LaNou



◆ Addison-Wesley

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

Cover image: T-flex/Shutterstock

Figures 7.1, 7.2: Microsoft

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

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

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

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

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

Visit us on the Web: informit.com/aw

Library of Congress Control Number: 2022943902

Copyright © 2023 Pearson Education, Inc.

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

ISBN-13: 978-0-13-791830-0

ISBN-10: 0-13-791830-5

ScoutAutomatedPrintCode

Pearson's Commitment to Diversity, Equity, and Inclusion

Pearson is dedicated to creating bias-free content that reflects the diversity of all learners. We embrace the many dimensions of diversity, including but not limited to race, ethnicity, gender, socioeconomic status, ability, age, sexual orientation, and religious or political beliefs.

Education is a powerful force for equity and change in our world. It has the potential to deliver opportunities that improve lives and enable economic mobility. As we work with authors to create content for every product and service, we acknowledge our responsibility to demonstrate inclusivity and incorporate diverse scholarship so that everyone can achieve their potential through learning. As the world's leading learning company, we have a duty to help drive change and live up to our purpose to help more people create a better life for themselves and to create a better world.

Our ambition is to purposefully contribute to a world where

- Everyone has an equitable and lifelong opportunity to succeed through learning.
- Our educational products and services are inclusive and represent the rich diversity of learners.
- Our educational content accurately reflects the histories and experiences of the learners we serve.
- Our educational content prompts deeper discussions with learners and motivates them to expand their own learning (and worldview).

While we work hard to present unbiased content, we want to hear from you about any concerns or needs with this Pearson product so that we can investigate and address them.

Please contact us with concerns about any potential bias at

<https://www.pearson.com/report-bias.html>.

This page intentionally left blank

For Rachel, Dylan, Leo, and Ringo.
—Mark Bates

For Karie, Logan, and Megan.
—Cory LaNou

This page intentionally left blank

Contents

Foreword xxi

Preface xxiii

Acknowledgments xxxi

About the Authors xxxv

1 Modules, Packages, and Dependencies 1

Modules	1
Toolchain	2
Initializing a Module	3
VCS and Modules	3
Packages	4
Naming Packages	5
Folders, Files, and Organization	6
Multiple Packages in a Folder	7
File Names	9
Package Organization	9
Importing Packages and Modules	13
Import Path	13
Using Imports	14
Resolving Import Name Clashes	14
Dependencies	16
Using a Dependency	17
Requiring with Go Get	18
The Go Sum File	20
Updating Dependencies	21
Semantic Versioning	22
Multiple Versions	23
Cyclical Imports	24
Summary	25

2	Go Language Basics	27
	Go Language Overview	27
	Static Typing	27
	Garbage Collection	28
	Compilation	28
	Keywords, Operators, and Delimiters	30
	Numbers	31
	Picking the Correct Numeric Type	32
	Overflow versus Wraparound	32
	Saturation	34
	Strings	35
	Interpreted String Literals	35
	Raw String Literals	36
	UTF-8	37
	Runes	38
	Iterating over UTF-8 Characters	38
	Variables	40
	Variable Declaration	40
	Variable Assignment	41
	Zero Values	42
	Nil	42
	Zero Values Cheat Sheet	43
	Variable Declaration and Initialization	44
	Assigning Multiple Values	45
	Unused Variables	46
	Constants	48
	Typed Constants	49
	Untyped Constants (Inferred Typing)	50
	Type Inference	50
	Naming Identifiers	51
	Naming Style	53
	Conflicting with Package Names	53
	Exporting through Capitalization	57
	Printing and Formatting	57
	Formatting Functions	57
	New Lines	60
	Multiple Arguments with Println	61
	Using Formatting Functions	61

Escape Sequences	61
Formatting Strings	63
Formatting Integers	63
Formatting Floats	66
Printing a Value's Type	67
Printing a Value	67
Printing a Value with More Detail	68
Printing a Value's Go-Syntax Representation	69
Using Formatting Verbs Incorrectly	70
Explicit Argument Indexes	71
Converting Strings to and from Numbers	72
Summary	75

3 Arrays, Slices, and Iteration 77

List Types: Arrays and Slices	77
Differences between Arrays and Slices	77
Spotting the Difference	79
Initializing Arrays and Slices	79
Array and Slice Zero Values	81
Indexing Arrays and Slices	81
Array and Slice Types	82
Array and Slice Type Definitions	83
Setting Array and Slice Values	85
Appending to Slices	85
Appending Slices to Slices	87
How Slices Work	90
Length and Capacity	91
Growing a Slice	92
Making a Slice	95
Make with Length and Capacity	96
Make and Append	97
What Happens When a Slice Grows	97
Slice Subsets	99
Mutating Slice Subsets	100
Copying Slices	101
Converting an Array to a Slice	102
Iteration	104
The <code>for</code> Loop	104
Iterating over Arrays and Slices	105

- The `range` Keyword 105
- Controlling Loops 107
- Do While Loop 108
- Summary 110

4 Maps and Control Structures 111

- Maps 111
 - Length and Capacity 112
 - Initializing Maps 113
 - Uninitialized Maps 114
 - Map Keys 114
 - Structs as Keys 115
 - Iterating Maps 116
 - Deleting Keys from a Map 118
 - Nonexistent Map Keys 120
 - Checking Map Key Existence 120
 - Exploiting Zero Value 121
 - Testing Presence Only 122
 - Maps and Complex Values 123
 - Copy on Insert 124
 - Listing Keys in a Map 126
- If Statements 129
 - The `else` Statement 130
 - The `else if` Statement 131
 - Assignment Scope 132
 - Logic and Math Operators 134
- `Switch` Statements 135
 - Default 138
 - Fallthrough 139
- Summary 140

5 Functions 141

- Function Definitions 141
 - Arguments 141
 - Arguments of the Same Type 142
 - Return Arguments 143
 - Multiple Return Arguments 144
 - Named Returns 145
 - First-Class Functions 147

Functions as Arguments	148
Closures	149
Anonymous Functions	150
Functions Accepting Arguments from Functions	151
Variadic Arguments	151
Variadic Argument Position	152
Expanding Slices	153
When to Use a Variadic Argument	154
Deferring Function Calls	156
Deferring with Multiple Returns	156
Defer Call Order	157
Deferred Calls and Panic	158
Defers and Exit/Fatal	158
Defer and Anonymous Functions	159
Defer and Scope	160
Init	162
Multiple Init Statements	163
Init Order	164
Using init Functions for Side Effects	164
Summary	166
6 Structs, Methods, and Pointers	167
Structs	167
Declaring New Types in Go	167
Defining a Struct	168
Initializing a Struct	169
Initialization without Field Names	170
Accessing Struct Fields	172
Struct Tags	173
Struct Tags for Encoding	173
Using Struct Tags	175
Methods	176
Differences Between Methods and Functions	177
Methods Expressions	177
Methods on Third-Party Types	178
Functions as Types	180
Methods on Functions	181

- No Inheritance 181
- Pointers 182
 - Pass by Value 183
 - Receiving a Pointer 184
 - Getting a Pointer 185
 - Passing with Pointers 186
 - Using Pointers 186
 - Value versus Pointer Receivers 188
 - New 189
 - Performance 190
- Nil Receivers 191
 - Nil Receivers Checks 192
- Summary 193

7 Testing 195

- Testing Basics 195
 - Naming 195
 - The `*testing.T` Type 196
 - Marking Test Failures 198
 - Using `t.Error` 198
 - Using `t.Fatal` (Recommended) 200
 - Crafting Good Test Failure Messages 201
- Code Coverage 202
 - Basic Code Coverage 202
 - Generating a Coverage Profile 203
 - The `go tool cover` Command 204
 - Generating an HTML Coverage Report 205
 - Editor Support 206
- Table Driven Testing 206
 - Anatomy of a Table Driven Test 207
 - Writing a Table Driven Test 208
 - Subtests 211
 - Anatomy of a Subtest 211
 - Writing Subtests 212
- Running Tests 213
 - Running Package Tests 213
 - Running Tests with Subpackages 213
 - Verbose Test Output 214
 - Logging in Tests 215

Short Tests	216
Running Package Tests in Parallel	217
Running Tests in Parallel	217
Running Specific Tests	218
Timing Out Tests	219
Failing Fast	220
Disabling Test Caching	221
Test Helpers	222
Defining Test Helpers	222
Marking a Function as a Helper	226
Cleaning Up a Helper	227
Summary	229

8 Interfaces 231

Concrete Types versus Interfaces	231
Explicit Interface Implementation	233
Implicit Interface Implementation	234
Before Interfaces	235
Using Interfaces	237
Implementing <code>io.Writer</code>	239
Multiple Interfaces	241
Asserting Interface Implementation	241
The Empty Interface	242
The <code>any</code> Keyword	242
The Problem with Empty Interfaces	243
Using an Empty Interface	243
Defining Interfaces	243
Defining a Model Interface	244
Implementing the Interface	247
Embedding Interfaces	249
Defining an Validatable Interface	250
Type Assertion	250
Asserting Assertion	251
Asserting Concrete Types	252
Assertions through Switch	252
Capturing Switch Type	253
Beware of Case Order	254
Using Assertions	255
Defining the Callback Interfaces	255

- Breaking It Down 256
- Summary 257

9 Errors 259

- Errors as Values 259
 - The `error` Interface 261
 - Handling Errors 262
 - Using Errors 263
- Panic 264
 - Raising a Panic 265
 - Recovering from a Panic 265
 - Capturing and Returning Panic Values 269
- Don't Panic 273
 - Checking for Nil 274
 - Maps 275
 - Pointers 278
 - Interfaces 280
 - Functions 283
 - Type Assertions 284
 - Array/Slice Indexes 287
- Custom Errors 289
 - Standard Errors 289
 - Defining Custom Errors 291
- Wrapping and Unwrapping Errors 294
 - Wrapping Errors 296
 - Unwrapping Errors 297
 - Unwrapping Custom Errors 298
 - To Wrap or Not To Wrap 301
- Errors As/Is 301
 - As 302
 - Is 304
- Stack Traces 307
- Summary 309

10 Generics 311

- What Are Generics? 311
 - The Problem with Interfaces 311
 - Type Constraints 315
 - Multiple Generic Types 317

Instantiating Generic Functions	320
Defining Constraints	321
Multiple Type Constraints	322
Underlying Type Constraints	324
The Constraints Package	326
Type Assertions	329
Mixing Method and Type Constraints	331
Generic Types	332
Summary	334

11 Channels 335

Concurrency and Parallelism	335
Concurrency Is Not Parallelism	336
Understanding Concurrency	336
Go's Concurrency Model	338
Goroutines	338
Goroutine Memory	339
The Go Scheduler	339
Work Sharing and Stealing	339
Don't Worry about the Scheduler	341
Goroutine Example	342
Communicating with Channels	342
What Are Channels?	343
Understanding Channel	
Blocking/Unblocking	343
Creating Channels	343
Sending and Receiving Values	344
A Simple Channel Example	345
Ranging over a Channel	348
Listening to Channels with <code>select</code>	348
Using <code>select</code> Statements	349
Channels Are Not Message Queues	349
Unidirectional Channels	351
Understanding Unidirectional Channels	352
Closing Channels	352
Detecting Closed Channels on Read	354
Zero Value on Closed Read	355
Closing an Already-Closed Channel	357

- Writing to a Closed Channel 358
- Buffered Channels 358
 - Basic Buffered Channel Example 359
 - Reading from Closed Buffered Channels 362
- Capturing System Signals with Channels 363
 - The `os/signals` Package 363
 - Implementing Graceful Shutdown 365
 - Listening for System Signals 367
 - Listening for Shutdown Confirmation 368
 - Timing Out a Nonresponsive Shutdown 370
- Summary 371

12 Context 373

- The Context Interface 374
 - `Context#Deadline` 374
 - `Context#Done` 375
 - `Context#Err` 376
 - `Context#Value` 376
 - Helper Functions 378
 - The Background Context 378
 - Default Implementations 379
- Context Rules 380
- Context Nodal Hierarchy 381
 - Understanding the Nodal Hierarchy 381
 - Wrapping with Context Values 382
 - Following the Context Nodes 382
- Context Values 384
 - Understanding Context Values 384
 - Key Resolution 386
- Problems with String Keys 388
 - Key Collisions 388
 - Custom String Key Types 390
- Securing Context Keys and Values 393
 - Securing by Not Exporting 394
- Cancellation Propagation with Contexts 396
 - Creating a Cancelable Context 397
 - Canceling a Context 398
 - Listening for Cancellation Confirmation 400
- Timeouts and Deadlines 405

Canceling at a Specific Time	405
Canceling after a Duration	407
Context Errors	408
Context Canceled Error	409
Context Deadline Exceeded Error	410
Listening for System Signals with Context	411
Testing Signals	413
Summary	416

13 Synchronization 419

Waiting for Goroutines with a WaitGroup	419
The Problem	419
Using a WaitGroup	421
The <code>Wait</code> Method	421
The <code>Add</code> Method	422
The <code>Done</code> Method	426
Wrapping Up Wait Groups	432
Error Management with Error Groups	433
The Problem	434
The Error Group	436
Listening for Error Group Cancellation	439
Wrapping Up Error Groups	440
Data Races	443
The Race Detector	445
Most, but Not All	446
Wrapping Up the Race Detector	447
Synchronizing Access with a Mutex	447
Locker	449
Using a Mutex	451
RWMutex	453
Improper Usage	455
Wrapping Up Read/Write Mutexes	459
Performing Tasks Only Once	459
The Problem	460
Implementing Once	462
Closing Channels with Once	463
Summary	466

14 Working with Files 467

- Directory Entries and File Information 467
 - Reading a Directory 468
 - The FileInfo Interface 471
 - Stating a File 472
- Walking Directories 473
- Skipping Directories and Files 477
 - Skipping Directories 478
- Creating Directories and Subdirectories 481
- File Path Helpers 484
 - Getting a File's Extension 485
 - Getting a File's Directory 485
 - Getting a File/Directory's Name 486
 - Using File Path Helpers 486
 - Checking the Error 488
- Creating Multiple Directories 489
- Creating Files 492
 - Truncation 495
- Fixing the Walk Tests 497
 - Creating the Files 498
- Appending to Files 500
- Reading Files 503
- Beware of Windows 505
- The FS Package 506
 - The FS Interface 508
 - The File Interface 509
- Using the FS Interface 510
 - File Paths 512
- Mocking a File System 513
 - Using MapFS 515
- Embedding Files 517
 - Using Embedded Files 518
 - Embedded Files in a Binary 520
 - Modifying Embedded Files 521
 - Embedding Files as a String or Byte Slice 522
- Summary 522

Index 523

Foreword

I feel excited, honored, and a bit shocked to be writing these words.

I've known Mark and Cory for more than a decade, and as the Go project has evolved beyond our 1.0 vision, I've really been hoping for someone to write a follow up to the excellent Go 1.0 books. I'm really grateful for the good fortune I've had to be so deeply involved in the Go project and community. What began with me typing the first few lines of code into Hugo,¹ Cobra,² and Viper³ over 10 years ago has exploded in ways I couldn't have dreamed of. It's been the privilege of a lifetime to lead the Go project alongside my partners Russ Cox and Sameer Ajmani and to work alongside the luminary programmers both on the Go team at Google and across the Go community.

I met Mark when we both spoke at the first Gotham Go.⁴ I was instantly drawn away from him. I found his on-stage persona to be a bit too much. Later that year, the organizers of GopherCon⁵ asked me to lead the lightning talk program and said they had the perfect partner for me. You guessed it—Mark. I learned that Mark's stage persona was indeed who he was off stage. I also learned that he was a loyal friend who would do anything for the Go community, whose humor and courage was boundless... limitless... maybe a bit too much. Mark and I have paired on many stages and projects over the past 10 years, and Mark injected a spirit of excitement into everything he did. He's become a dear friend, and I'm grateful to have been on so many adventures with him.

I met Cory at the second GopherCon and instantly was drawn to him. Cory is a natural teacher—captivating and empathetic. He cares deeply about the Go community and ensuring that Go is accessible to all, especially the folks new to Go. I've worked with him on a variety of community efforts, and I've always come away impressed with the depth he shows in his knowledge about Go and the learning experience.

Mark and Cory have been working together for many years now as the dominant training duo for Go under the name Gopher Guides. Together they have produced excellent training programs for clients that include many notable brands found on the Fortune 500. They have expertise in both Go and in empathic learning, forged over thousands of hours of classroom-style instruction. They are the perfect pair to author what is destined to become “The Go Book” for Go's second phase (Go with Modules & Generics).

This book leans on Mark and Cory's years of experience in the field and in the Go community to take a very grassroots approach to learning, and it is the guide for

1. <https://gohugo.io/>

2. <https://pkg.go.dev/github.com/spf13/cobra>

3. <https://pkg.go.dev/github.com/spf13/viper>

4. <https://gothamgo.com/>

5. <https://www.gophercon.com/>

programmers to become gophers. It also leverages their practical experience writing Go libraries and applications to present to the reader pragmatic solutions and simple explanations.

This book guides you like an old friend would, telling you the technical approaches to things but also relaying the cultural norms and idioms. It answers questions that the majority of books don't even think of because Mark and Cory have heard these questions asked by real people in a classroom so many times. Through reading and applying what is in this book, more than any other of its kind, you will progress from becoming a programmer to being a programmer who writes Go to being a gopher.

I am excited for the journey you are about to take with this book. I'm sure that you will sense the same excitement and empathy through these pages that I have experienced through all the adventures I've had with Mark and Cory. My hope for you is that as you learn Go, you will, as I did when I first discovered Go over a decade ago, fall in love with programming again.

—*Steve Francia*
@spf13

Preface

Over the years, we have been fortunate to have trained thousands of developers in Go. It has been an incredible experience for us. We are developers who care about developers, so to be able to help developers grow is a great honor for us. This book is a culmination of that experience.

With nearly a half-century of experience between us, we both love Go. Go is fast, efficient, and fairly straightforward. We both believe that Go is a great first language to learn.

Unlike previous languages we have both used, there is relatively little “magic” in the language. Developers are rarely left scratching their heads as to where a function or type came from. Although the language does offer tools, such as the `reflect`¹ package, those tools are not used often. As part of this philosophy, the general consensus in the community is to favor using the standard library over third-party libraries. Again, this stems from experience with languages where bloated dependencies caused all sorts of problems. Also, early in the language’s history, there was no package manager, so managing dependencies was a challenge.

The Go compiler is another reason why Go has become so popular. The speed of the compiler is a huge factor in the development of Go. Fast compilation times mean a faster feedback cycle for developers. Large Go applications can be compiled in seconds. Small Go applications, when using `go run`, compile and execute so quickly that it makes Go feel like a scripting language. This fast compilation and execution time extends to the running of tests as well. It is not uncommon for a Go developer to run the entire test suite every time they save a file. This is not something that can be done with other languages.

The Go compiler, the language’s type system, and its fast compiler mean that a lot of common bugs are very quickly caught. This is a great advantage for Go developers. We are able to focus on business logic, knowing that the compiler will catch any stupid mistakes we might make, such as using a variable before it is defined or forgetting to use the `range` keyword in a loop.

For example, in Listing P.1, we are doing just that. You can see from the output that the compiler caught the error and did so very quickly.

1. <https://pkg.go.dev/reflect>

Listing P.1 Quick Feedback Cycle Thanks to the Go Compiler

```
package main

import (
    "fmt"
    "os"
    "os/exec"
    "time"
)

func main() {
    letters := []string{"a", "b", "c"}

    for _, l := letters {
        fmt.Println(l)
    }
}

$ go run .

# demo
./main.go:13:11: syntax error: cannot use _, l := letters as value

Duration:    64.341084ms
Go Version:  go1.19
```

We discuss more about looping in Chapter 3 and more about how the Go compiler and runtime work as we progress throughout the book.

Finally, Go has a great concurrency story. The language was built from the ground with concurrency in mind. This means you don't need to worry about threading, spawning processes, or anything like that. You can write concurrent code quickly and easily.

In the code snippet in Listing P.2, we are using goroutines, channels, and synchronization types, such as `sync.WaitGroup`² to create an application that properly handles concurrency.

Listing P.2 Go Concurrency in Action

```
func main() {
    // create a channel of type int
    ch := make(chan int)
```

2. <https://pkg.go.dev/sync#WaitGroup>

```
// create a wait group to track
// the goroutines
var wg sync.WaitGroup

// create some goroutines to
// listen for messages on the
// channel
for i := 0; i < 4; i++ {

    // increment the wait group
    wg.Add(1)

    // create a goroutine
    // to call the sayHello function
    go func(id int) {

        // call the sayHello function
        sayHello(id, ch)

        // decrement the wait group
        // when the sayHello function
        // exits
        wg.Done()
    }(i + 1)
}

// send messages on the channel
for i := 0; i < 10; i++ {
    ch <- i
}

// close the channel to signal
// the goroutines to exit
close(ch)

// wait for all goroutines to finish
wg.Wait()
}

func sayHello(id int, ch chan int) {

    // listen for messages on the channel
    // loop exits when the channel is closed
    for i := range ch {
```

```
    fmt.Printf("Hello %d from goroutine %d\n", i, id)

    // simulate a long-running task
    sleep()
}
}
```

```
$ go run .
```

```
Hello 2 from goroutine 4
Hello 1 from goroutine 2
Hello 3 from goroutine 3
Hello 0 from goroutine 1
Hello 4 from goroutine 1
Hello 5 from goroutine 2
Hello 6 from goroutine 1
Hello 7 from goroutine 1
Hello 8 from goroutine 4
Hello 9 from goroutine 3
```

```
Go Version: go1.19
```

You can read more about Go's concurrency model in Chapters, 11, 12, and 13.

How to Read This Book

We have laid this book out with the idea that it is meant to be read from start to finish. Each chapter, and each example, build on previous examples. While those with some experience with Go might be tempted to jump straight to later chapters, we really recommend reading the whole book.

We have gone through great effort to try to make sure that examples and chapters don't contain any new information that we haven't covered up to that point. In Chapter 9, for example, we cover errors in Go. To understand how errors work, you first need to learn about interfaces, which we cover in Chapter 8. Occasionally, we do have to break this rule, though. In those instances, we try our best to explain the concept or point to the chapter where it is explained.

About the Examples

As mentioned, we spent considerable time designing the examples in this book and our training materials to be as clear and concise as possible. We also wanted to make sure that all of the code, output, and documentation were accurate.

To ensure that all of these things are true, we used a system designed give us this accuracy. Knowing that every code snippet is from a real file, every command is actually executed, every output is exact, and every piece of documentation is up to date means we can be confident in the materials we present.

In Listing P.3, you see an example of what one of these documents looks like.

Listing P.3 Sample of How the Book Was Written

Commands and Documentation

In [exit](#), we see a snippet of code.

This will execute the command `'go run .'` in the `'src/bad'` directory. This command is expected to exit with a code of `'1'`. The output of the command is automatically captured and inserted into the document.

```
<figure id="exit" type="listing">
<code src="module.md#exit" esc></code>
<figcaption>Handling a non-successful exit code.</figcaption>
</figure>
```

In [panic](#), we can see the output of the non-successful program.

Notice, at the top of the output we see the command used that was executed, `'$ go run .'`. At the bottom we can see that the program was executed with, Go version `'1.19'`.

```
<figure id="panic">
<go src="src/bad" run="." exit="1"></go>
<figcaption>Output of a non-successful program.</figcaption>
</figure>
```

While this book was written with Go version `'1.18'`, all of the final examples, and documentation, were run with Go version `'1.19'`.

In [appendf.doc](#), we can see an example of documentation that is being inserted into the document. In this case, we are inserting the documentation for the new Go `'1.19'` function, `'fmt.Appendf'`. Again, just like the command output in [panic](#), the command used to get the documentation and the Go version is shown at the bottom of the output.

(continued)

```
<figure id="appendf.doc">
<go doc="fmt.Appendf"></go>
<figcaption>The 'fmt.Appendf' function documentation.</figcaption>
</figure>
```

The output of Listing P.3 is included in the next section.

Commands and Documentation

Listing P.4 shows a snippet of code.

It executes the command `go run .` in the `src/bad` directory. This command is expected to exit with a code of `1`. The output of the command is automatically captured and inserted into the document.

Listing P.4 Handling a Nonsuccessful Exit Code

```
<go src="src/bad" run="." exit="1"></go>
```

In Listing P.5, you can see the output of the nonsuccessful program. Notice that the top of the output, you see the command used that was executed: `$ go run .`. At the bottom, you can see that the program was executed with Go version `1.19`.

Listing P.5 Output of a Nonsuccessful Program

```
$ go run .

panic: Hello, World!

goroutine 1 [running]:
main.main()
    ./main.go:5 +0x2c
exit status 2
```

```
Go Version: go1.19
```

While this book was written with Go version `1.18`, all of the final examples and documentation were run with Go version `1.19`.

In Listing P.6, you can see an example of documentation that is being inserted into the document. In this case, we are inserting the documentation for the new Go `1.19` function, `fmt.Appendf`. Again, just like the command output in Listing P.5, the command used to get the documentation, `$ go do fmt.Appendf`, is shown at the top of the output, and the Go version is shown at the bottom of the output.

Listing P.6 The `fmt.Appendf` Function Documentation

```
$ go doc fmt.Appendf

package fmt // import "fmt"

func Appendf(b []byte, format string, a ...any) []byte
    Appendf formats according to a format specifier, appends the result to the
    byte slice, and returns the updated slice.

Go Version: go1.19
```

Summary

Contained within these pages are the concepts, types, packages, idioms, and other features of Go that we believe to be the most fundamental to understanding the language.

Go, its ecosystem, and its community are large, complex, and vibrant. Although we wish we could have included more in this book, it would've been impossible to print, and we would have never finished it.

Our goal, by the end of this book, is to help you be a knowledgeable Go programmer—one who not only feels confident using the language but understands how to write idiomatic Go code, tests, and is a good community member.

Finally, dear reader, thank you for your support. We consider it a privilege to be able to share this book and our knowledge with you. We love Go, and we hope you do, too.

—Mark Bates and Cory LaNou
August 2022

Register your copy of *Go Fundamentals* on the InformIT site for convenient access to updates and/or corrections as they become available. To start the registration process, go to informit.com/register and log in or create an account. Enter the product ISBN (9780137918300) and click Submit. Look on the Registered Products tab for an Access Bonus Content link next to this product, and follow that link to access any available bonus materials. If you would like to be notified of exclusive offers on new editions and updates, please check the box to receive email from us.

This page intentionally left blank

Acknowledgments

Authoring a book is incredibly hard work, whether you are self-publishing or working with a great publisher like Addison-Wesley. It requires long hours of writing, creating just the right sample code, and refactoring. This is usually where the self-publishing author stops. When working with a publisher, however, this is only the start of a new phase in the process.

After a book is written and submitted, a team of people jump in to try and make the book the best it can be. There are technical reviewers who are there to make sure the code and the discussion of the code are correct. There are copy editors who drill into the minutiae of grammar, spelling, punctuation, and other stylistic issues. There is the production team who is responsible for the final delivery of the book. There is the cover designer, the proofreader, the indexer, and more. As the adage says, “It takes a village.”

Gopher Guides would like to thank Kim Spenceley, our editor, and Debra Williams-Cauley, our managing editor, for all of their incredible support since we first started this journey. Their optimism and excitement has kept us going through this process. Without Kim and Debra, there would be no book. Thank you two so very much.

In addition to Kim and Debra, Gopher Guides would like to thank the following people at Addison-Wesley for their help in the process of making this book a reality:

Copy editor and project manager: Charlotte Kughen

Production manager: Sandra Schroeder

Production editor: Julie Nahil

Cover designer: Chuti Prasertsith

Development editor: Chris Zahn

Indexer: Cheryl Lenser

Proofreader: Sarah Kearns

The rest of the village at Addison-Wesley

The Go community has been a huge supporter of Gopher Guides over the years. We would like to thank everyone on the Go team for their dedication and hard work in creating a language that is loved by so many. Thank you to all those who have committed to the Go project over the years. Again, it takes a village.

In particular, Gopher guides would like to thank the following Go community members who have helped us along the way.

Steve Francia	Ron Evans
Ashley Willis (McNamara) (who designed the Gopher Guides logo)	Bill Kennedy
Mat Ryer	Brian Ketelsen
Carmen Andoh	Bryan Liles
Johnny Boursiquot	Ben Johnson
Tim Raymond	Dave Cheney
Antonio Pagano	Matt Aimonetti
Renee French (for creating the Gopher logo)	Francesc Campoy

From Mark Bates

In addition to everyone that Cory and I have thanked already, I would like to make a few acknowledgments of my own.

First, I would like to thank my business partner and friend, Cory LaNou. Cory and I knew each other well in the community, but we really got to know each other and became good friends on a trip to GopherCon India in Bengaluru (Bangalore), India, and in Dubai, United Arab Emirates. We both spoke at the conferences, which were great, but it was the extra curriculum that truly bonded us. We rode camels, drove dune buggies, and went to a desert oasis. It was a great experience.

About a year later, in 2017, Cory and I decided that we wanted to start a company to help people learn Go. With that, Gopher Guides was born. I have never worked anywhere for longer than two and a half years. Gopher Guides, at the time of writing, has been going for more than five years. That is more than double the time I've worked at any job. The reason for this is simple: Cory LaNou is an excellent partner and a great friend. I have never felt as supported and as understood while working at any other company as I do working with Cory.

Cory, thank you for your patience, understanding, and support. I couldn't be prouder of what we have built together—this great company and our enduring friendship.

Next, I want to thank Kim Spenceley and Debra Williams-Cauley. I have known both of them for a long time. Debra was the editor on my first two books for Addison-Wesley in 2009 and 2012. Kim was the publishing coordinator of the first book. Both of these strong women have been of great support to me over the years. I always feel lucky that I was able to find a publishing company like Addison-Wesley—with editors like Debra and Kim—that is so focused on author success. Thank you both for helping me to first realize my dream of having a book published, and then for encouraging me to continue writing.

Finally, I would like to thank my family. There is no better partner, friend, and supporter than my wife, Rachel. She is an actual wonder woman. The saying “behind every good man is a great woman” is about her. She is a ball of energy, enthusiasm, and spirit. She has a goal of running a half marathon in every U.S. state, and she's more than

halfway there. I get winded driving 13.1 miles. She is the morning person to my night owl. Rachel is a successful business woman who has spent her career being an advocate for women in the workplace. She is the mentor that she never had to a generation of women in business. I know several of these women, and to hear them talk about Rachel with such respect and appreciation is heart-warming.

Lastly, in addition to all of her support, love, and dedication over the years, Rachel has given me the greatest gift of all; our two sons, Dylan and Leo, are wonderful people. They are kind, caring, and honest. They are both very smart, talented, and hilarious people who continue to amaze me and make me glow with pride. Each night, we eat dinner together and then settle in, as a family, to watch TV. Our dog Ringo curls up at the end of the sofa, and we all enjoy each other's company. It's at those moments I feel truly thankful and happy to be surrounded by their love.

Thank you Rachel, Dylan, and Leo (and Ringo—the dog, not the Beatle—though he's cool, too!) for all of your support and love. I love you all to the very depths of my heart.

From Cory LaNou

First, and most obvious, my thanks go to my business partner, Mark Bates. Mark and I met through the Go community and have become not just great partners in our business but also great friends. We both share a passion for Go and training, which is why I believe we have been so successful together.

A big thanks to Levi Cook for introducing me to the Go language in March 2012 and encouraging me to start one of the first Go meetups in the world in Denver, Colorado. Levi not only introduced me to Go but also to open source and, most important, VIM!

I also want to recognize some of the brightest and most talented developers I have ever worked with on the InfluxDB project. Ben Johnson for having an ability to name the impossible things. Jason Wilder for opening my mind to taking the most complex of issues and breaking them into their smallest components. Philip O'Toole for challenging every new concept I introduced to the Go code base. David Norton for being the nicest person I've ever worked with (and giving the best SQL presentation ever!). Joe LeGasse for never hesitating to pair program with me. And all the other amazing people I worked with on the InfluxDB team.

When I moved back to Wisconsin to be around all my family, I knew I would miss Denver's rich technology community. Thankfully, Doug Rhoten was heading up the Chippewa Valley Developers meetup and made me feel incredibly welcome. Thank you, Doug, for working so hard for the local technology community!

And lastly, but certainly not least, my family! My wife, Karie, has been nothing but supportive, even when I have done the craziest of business ventures. Without her support and belief in me, I would not be where I am today. My son, Logan, and daughter, Megan, who thank me for all the long hours I put in so that they can have bacon for breakfast!

This page intentionally left blank

About the Authors

Mark Bates is cofounder and instructor at Gopher Guides, the industry leader for Go training, consulting, and conference workshops. After graduating with a degree in music from the Liverpool Institute for Performing Arts in 1999, Mark joined the original dot-com boom as a software engineer and began his career as a technologist. Since then, Mark has been fortunate enough to work with some of the world's largest and most innovative companies, including Apple, Uber, and Visa.

When he first discovered Go in the summer of 2013, Mark was immediately drawn to the language and its ecosystem. In 2014, Mark attended the first GopherCon, where he met Cory LaNou. For seven years, Mark hosted the GopherCon lightning talks and has been proud to introduce hundreds of new speakers to the community. Mark has spoken at conferences around the world and is a regular on the *Go Time* podcast.

When not coding or writing about coding, Mark enjoys spending time with his family, traveling, and recording music. Mark is currently completing a master's degree in music production from Berklee College of Music in Boston, Massachusetts.

Cory LaNou is a full stack technologist who has specialized in start-ups for the last 20 years. Cory has started several technology companies over the years, including companies such as Pulsity, a web consulting firm that operated in the late 1990s. From that company, Local Launch (an Internet marketing technology company) was formed and later sold to RH Donnelly in 2006.

Cory has deep experience in the Go world, having worked on the Core engineering team that contributes to InfluxDB, which is a highly scalable, distributed time series database written in Go. Cory has worked on several other Go projects with a focus on profiling and optimizing applications, using advanced profiling concepts such as torch graphs and tracing profiles.

Cory has deep ties to the Go community and started Denver Gophers, one of the very first Go meetups in the world, and he assisted in the very first GopherCon. Cory has created and led numerous Go workshops and training courses and has several published online articles related to Go. He continues to help organize and lead several community technology meetups and mentors new developers. Cory is also a partner at Gopher Guides, the industry leader for Go training and conferences.

When Cory isn't working on writing new Go training material, he volunteers his time to help local entrepreneurs start up their businesses by offering free services and business advice—such as how to incorporate their business and get their web presence set up—and giving them the moral support they need to make their dreams a reality.

This page intentionally left blank

Synchronization

In Chapter 11, we explained how to use channels for passing data between goroutines. Then in Chapter 12, we discussed how to use the `context`¹ package to manage the cancellation of goroutines. In this chapter, we cover the final part of concurrent programming: synchronization.

We show you how to wait for a number of goroutines to finish. We explain race conditions,² how to find them using Go's `-race`³ flag, and how to fix them with `sync.Mutex`⁴ and `sync.RWMutex`.⁵

Finally, we discuss how to use `sync.Once` to ensure a function is only executed one time.

Waiting for Goroutines with a WaitGroup

Often, you might want to wait for a number of goroutines to finish before you continue your program. For example, you might want to spawn a number of goroutines to create a number of thumbnails of different sizes and wait for them all to complete before you continue.

The Problem

Consider Listing 13.1. We launch 5 new goroutines, each of which creates a thumbnail of a different size. We then wait for all of them to complete.

Listing 13.1 Launching Multiple Goroutines to Complete One Task

```
func Test_ThumbnailGenerator(t *testing.T) {
    t.Parallel()

    // image that we need thumbnails for
    const image = "foo.png"
```

(continued)

1. <https://pkg.go.dev/context>
2. https://en.wikipedia.org/wiki/Race_condition
3. https://golang.org/doc/articles/race_detector
4. <https://pkg.go.dev/sync#Mutex>
5. <https://pkg.go.dev/sync#RWMutex>

```

// start 5 goroutines to generate thumbnails
for i := 0; i < 5; i++ {

    // start a new goroutine for each thumbnail
    go generateThumbnail(image, i+1)

}

fmt.Println("Waiting for thumbnails to be generated")
}

```

The `generateThumbnail` function, Listing 13.2, generates a thumbnail of the specified size. In this example, we sleep one millisecond per “size” of thumbnail to simulate the time it takes to generate the thumbnail. For example, if we call `generateThumbnail("foo.png", 200)`, we sleep 200 milliseconds before returning.

Listing 13.2 A Test Exiting before All Goroutines Have Finished

```

func generateThumbnail(image string, size int) {

    // thumbnail to be generated
    thumb := fmt.Sprintf("%s@%dx.png", image, size)

    fmt.Println("Generating thumbnail:", thumb)

    // wait for the thumbnail to be ready
    time.Sleep(time.Millisecond * time.Duration(size))

    fmt.Println("Finished generating thumbnail:", thumb)
}

```

```

$ go test -v

=== RUN   Test_ThumbnailGenerator
=== PAUSE Test_ThumbnailGenerator
=== CONT Test_ThumbnailGenerator
Waiting for thumbnails to be generated
--- PASS: Test_ThumbnailGenerator (0.00s)
PASS
ok      demo    0.408s

```

```
Go Version: go1.19
```

As you can see from the test output in Listing 13.2, the test exits before the thumbnails are generated.

Our tests exit prematurely because we have not provided any mechanics to ensure that we wait for all of the thumbnail goroutines to finish before we continue.

Using a WaitGroup

To help us solve this problem, we can use a `sync.WaitGroup`,⁶ Listing 13.3, to track how many goroutines are still running and notify us when they have all finished.

Listing 13.3 The `sync.WaitGroup` Type

```
$ go doc -short sync.WaitGroup

type WaitGroup struct {
    // Has unexported fields.
}

A WaitGroup waits for a collection of goroutines to finish. The main
↳goroutine calls Add to set the number of goroutines to wait for. Then
↳each of the goroutines runs and calls Done when finished. At the same
↳time, Wait can be used to block until all goroutines have finished.

A WaitGroup must not be copied after first use.

func (wg *WaitGroup) Add(delta int)
func (wg *WaitGroup) Done()
func (wg *WaitGroup) Wait()

Go Version: go1.19
```

The principle is simple: We create a `sync.WaitGroup` and use the `sync.WaitGroup.Add`⁷ method to add to the `sync.WaitGroup` for each goroutine we want to wait for. When we want to wait for all of the goroutines to finish, we call the `sync.WaitGroup.Wait`⁸ method. When a goroutine finishes, it calls the `sync.WaitGroup.Done`⁹ method to indicate that the goroutine is finished.

The Wait Method

As the name suggests, a `sync.WaitGroup` is about waiting for a group of tasks, or goroutines, to finish. To do this, we need a way of blocking until all of the tasks have finished. The `sync.WaitGroup.Wait` method in Listing 13.4 does exactly that.

6. <https://pkg.go.dev/sync#WaitGroup>

7. <https://pkg.go.dev/sync#WaitGroup.Add>

8. <https://pkg.go.dev/sync#WaitGroup.Wait>

9. <https://pkg.go.dev/sync#WaitGroup.Done>

The `sync.WaitGroup.Wait` method blocks until its internal counter is zero. When the counter is zero, it means that all of the tasks have finished, and we can unblock and continue.

Listing 13.4 The `sync.WaitGroup.Wait` Method

```
$ go doc sync.WaitGroup.Wait

package sync // import "sync"

func (wg *WaitGroup) Wait()
    Wait blocks until the WaitGroup counter is zero.
```

Go Version: go1.19

The Add Method

For a `sync.WaitGroup` to know how many goroutines it needs to wait for, we need to add them to the `sync.WaitGroup` using the `sync.WaitGroup.Add` method, Listing 13.5.

Listing 13.5 The `sync.WaitGroup.Add` Method

```
$ go doc sync.WaitGroup.Add

package sync // import "sync"

func (wg *WaitGroup) Add(delta int)
    Add adds delta, which may be negative, to the WaitGroup counter.
    ➤If the counter becomes zero, all goroutines blocked on Wait are released.
    ➤If the counter goes negative, Add panics.
```

Note that calls with a positive delta that occur when the counter is

- zero must happen before a `Wait`. Calls with a negative delta, or calls
- with a positive delta that start when the counter is greater than zero,
- may happen at any time. Typically this means the calls to `Add` should
- execute before the statement creating the goroutine or other event to be
- waited for. If a `WaitGroup` is reused to wait for several independent
- sets of events, new `Add` calls must happen after all previous `Wait` calls
- have returned. See the `WaitGroup` example.

Go Version: go1.19

The `sync.WaitGroup.Add` method takes a single integer argument, which is the number of goroutines to wait for. There are, however, some caveats to be aware of.

Adding a Positive Number

The `sync.WaitGroup.Add` method accepts an `int` argument, which is the number of goroutines to wait for. If we pass a positive number, the `sync.WaitGroup.Add` method adds that number of goroutines to the `sync.WaitGroup`.

As you can see from the test output in Listing 13.6, the `sync.WaitGroup.Wait` method blocks until the internal counter of the `sync.WaitGroup` reaches zero.

Listing 13.6 Adding a Positive Number of Goroutines

```
func Test_WaitGroup_Add_Positive(t *testing.T) {
    t.Parallel()

    var completed bool

    // create a new waitgroup (count: 0)
    var wg sync.WaitGroup

    // add one to the waitgroup (count: 1)
    wg.Add(1)

    // launch a goroutine to call the Done() method
    go func(wg *sync.WaitGroup) {

        // sleep for a bit
        time.Sleep(time.Millisecond * 10)

        fmt.Println("done with waitgroup")

        completed = true

        // call the Done() method to decrement
        // the waitgroup counter (count: 0)
        wg.Done()
    }(&wg)

    fmt.Println("waiting for waitgroup to unblock")

    // wait for the waitgroup to unblock (count: 1)
    wg.Wait()
    // (count: 0)
```

(continued)

```

    fmt.Println("waitgroup is unblocked")

    if !completed {
        t.Fatal("waitgroup is not completed")
    }
}

```

```

$ go test -v -run Positive

=== RUN   Test_WaitGroup_Add_Positive
=== PAUSE Test_WaitGroup_Add_Positive
=== CONT Test_WaitGroup_Add_Positive
waiting for waitgroup to unblock
done with waitgroup
waitgroup is unblocked
--- PASS: Test_WaitGroup_Add_Positive (0.01s)
PASS
ok      demo    0.351s

```

```

Go Version: go1.19

```

Adding a Zero Number

It is legal to call the `sync.WaitGroup.Add` method with a zero number, 0, Listing 13.7. In this case, the `sync.WaitGroup.Add` method does nothing. The call becomes a no-op.

Listing 13.7 Adding a Zero Number of Goroutines

```

func Test_WaitGroup_Add_Zero(t *testing.T) {
    t.Parallel()

    // create a new waitgroup (count: 0)
    var wg sync.WaitGroup

    // add 0 to the waitgroup (count: 0)
    wg.Add(0)
    // (count: 0)

    fmt.Println("waiting for waitgroup to unblock")

    // wait for the waitgroup to unblock (count: 0)
    // will not block since the counter is already 0

```

(continued)

```

wg.Wait()
// (count: 0)

fmt.Println("waitgroup is unblocked")
}

```

```

$ go test -v -run Zero

=== RUN   Test_WaitGroup_Add_Zero
=== PAUSE Test_WaitGroup_Add_Zero
=== CONT Test_WaitGroup_Add_Zero
waiting for waitgroup to unblock
waitgroup is unblocked
--- PASS: Test_WaitGroup_Add_Zero (0.00s)
PASS
ok      demo    0.166s

```

```

Go Version: go1.19

```

As you can see from the test output in Listing 13.7, the `sync.WaitGroup.Wait` method unblocked immediately because its internal counter is already zero.

Adding a Negative Number

When calling the `sync.WaitGroup.Add` method with a negative number, the `sync.WaitGroup.Add` method panics.

As you can see from the test output in Listing 13.8, the `sync.WaitGroup.Wait` method was never reached because the `sync.WaitGroup.Add` method panicked when we tried to add a negative number of goroutines.

Listing 13.8 Adding a Negative Number of Goroutines

```

func Test_WaitGroup_Add_Negative(t *testing.T) {
    t.Parallel()

    // create a new waitgroup (count: 0)
    var wg sync.WaitGroup

    // use an anonymous function to trap the panic
    // so we can properly mark the test as a failure
    func() {

        // defer a function to catch the panic
        defer func() {

```

(continued)

```

        // recover the panic
        if r := recover(); r != nil {
            // mark the test as a failure
            t.Fatal(r)
        }
    }()

    // add a negative number to the waitgroup
    // this will panic since the counter cannot be negative
    wg.Add(-1)

    fmt.Println("waiting for waitgroup to unblock")

    // this will never be reached
    wg.Wait()

    fmt.Println("waitgroup is unblocked")
}()
}

```

```

$ go test -v -run Negative

=== RUN   Test_WaitGroup_Add_Negative
=== PAUSE Test_WaitGroup_Add_Negative
=== CONT  Test_WaitGroup_Add_Negative
    add_test.go:92: sync: negative WaitGroup counter
--- FAIL: Test_WaitGroup_Add_Negative (0.00s)
FAIL
exit status 1
FAIL    demo    0.753s

```

```

Go Version: go1.19

```

The Done Method

Once we increase that counter by calling the `sync.WaitGroup.Add` method, the `sync.WaitGroup.Wait` method blocks until we decrement the counter as we finish with each goroutine.

For each item we add to the `sync.WaitGroup` with the `sync.WaitGroup.Add` method, we need to call the `sync.WaitGroup.Done` method, Listing 13.9, to indicate that the goroutine is finished.

Listing 13.9 The `sync.WaitGroup.Done` method

```
$ go doc sync.WaitGroup.Done

package sync // import "sync"

func (wg *WaitGroup) Done()
    Done decrements the WaitGroup counter by one.
```

```
Go Version: go1.19
```

Consider Listing 13.10, which creates `N` goroutines and adds `N` to the `sync.WaitGroup` using the `sync.WaitGroup.Add` method. Each goroutine calls the `sync.WaitGroup.Done` method after it finishes. We then use the `sync.WaitGroup.Wait` method to wait for all of the goroutines to finish.

Listing 13.10 Testing the `sync.WaitGroup.Done` Method

```
func Test_WaitGroup_Done(t *testing.T) {
    t.Parallel()

    const N = 5

    // create a new waitgroup (count: 0)
    var wg sync.WaitGroup

    // add 5 to the waitgroup (count: 5)
    wg.Add(N)

    for i := 0; i < N; i++ {

        // launch a goroutine that will call the
        // waitgroup's Done method when it finishes
        go func(i int) {

            // sleep briefly
            time.Sleep(time.Millisecond * time.Duration(i))

            fmt.Println("decrementing waiting by 1")

            // call the waitgroup's Done method
            // (count: count - 1)
            wg.Done()
        }()
    }
}
```

(continued)

```

    }(i + 1)
}

fmt.Println("waiting for waitgroup to unblock")

wg.Wait()

fmt.Println("waitgroup is unblocked")
}

```

```

$ go test -v -timeout 1s

=== RUN   Test_WaitGroup_Done
=== PAUSE Test_WaitGroup_Done
=== CONT  Test_WaitGroup_Done
waiting for waitgroup to unblock
decremeting waiting by 1
decremeting waiting by 1
decremeting waiting by 1
decremeting waiting by 1
decremeting waiting by 1
decremeting waiting by 1
waitgroup is unblocked
--- PASS: Test_WaitGroup_Done (0.01s)
PASS
ok      demo    0.384s

```

```

Go Version: go1.19

```

As we can see from the test output, Listing 13.10, the `sync.WaitGroup.Wait` method unblocked after all of the goroutines finished.

Improper Usage

You must call `sync.WaitGroup.Done` exactly once for each number of items you add with `sync.WaitGroup.Add`.

If you don't call `sync.WaitGroup.Done` exactly once for each item you add with `sync.WaitGroup.Add`, the `sync.WaitGroup.Wait` method will block forever, which causes a deadlock and crashes your program, as shown in Listing 13.11.

Listing 13.11 Decrementing a `sync.WaitGroup` with the `sync.WaitGroup.Done` Method

```
func Test_WaitGroup_Done(t *testing.T) {
    t.Parallel()

    const N = 5

    // create a new waitgroup (count: 0)
    var wg sync.WaitGroup

    // add 5 to the waitgroup (count: 5)
    wg.Add(N)

    for i := 0; i < N; i++ {

        // launch a goroutine that will call the
        // waitgroup's Done method when it finishes
        go func(i int) {

            // sleep briefly
            time.Sleep(time.Millisecond * time.Duration(i))

            fmt.Println("finished")

            // exiting with calling the Done method
            // (count: count)
        }(i + 1)
    }

    fmt.Println("waiting for waitgroup to unblock")

    // this will never unblock
    // because the goroutines never call Done
    // and the application will deadlock and panic
    wg.Wait()

    fmt.Println("waitgroup is unblocked")
}
```

```
$ go test -v -timeout 1s
=== RUN   Test_WaitGroup_Done
=== PAUSE Test_WaitGroup_Done
=== CONT Test_WaitGroup_Done
```

(continued)


```

waiting for waitgroup to unblock
finished
finished
finished
finished
finished
panic: test timed out after 1s

goroutine 19 [running]:
testing.(*M).startAlarm.func1()
    /usr/local/go/src/testing/testing.go:2029 +0x8c
created by time.goFunc
    /usr/local/go/src/time/sleep.go:176 +0x3c

goroutine 1 [chan receive]:
testing.tRunner.func1()
    /usr/local/go/src/testing/testing.go:1405 +0x45c
testing.tRunner(0x140001361a0, 0x1400010fcb8)
    /usr/local/go/src/testing/testing.go:1445 +0x14c
testing.runTests(0x1400001e280?, {0x101045ea0, 0x1, 0x1},
    ↪{0x6e00000000000000?, 0x100e71218?, 0x10104e640?})
    /usr/local/go/src/testing/testing.go:1837 +0x3f0
testing.(*M).Run(0x1400001e280)
    /usr/local/go/src/testing/testing.go:1719 +0x500
main.main()
    _testmain.go:47 +0x1d0

goroutine 4 [semacquire]:
sync.runtime_Semacquire(0x0?)
    /usr/local/go/src/runtime/sema.go:56 +0x2c
sync.(*WaitGroup).Wait(0x14000012140)
    /usr/local/go/src/sync/waitgroup.go:136 +0x88
demo.Test_WaitGroup_Done(0x0?)
    ./done_test.go:43 0xd0
testing.tRunner(0x14000136340, 0x100fa1580)
    /usr/local/go/src/testing/testing.go:1439 +0x110
created by testing.(*T).Run
    /usr/local/go/src/testing/testing.go:1486 +0x300
exit status 2
FAIL    demo    1.225s

```

Go Version: go1.19

If you call `sync.WaitGroup.Done` more than the number of items you added with `sync.WaitGroup.Add`, the `sync.WaitGroup.Done` method panics, Listing 13.12. The result is the same as if you called `sync.WaitGroup.Add` with a negative number.

Listing 13.12 Panicking from Decrementing `sync.WaitGroup` Too Many Times

```
func Test_WaitGroup_Done(t *testing.T) {
    t.Parallel()

    func() {
        // defer a function to catch the panic
        defer func() {

            // recover the panic
            if r := recover(); r != nil {
                // mark the test as a failure
                t.Fatal(r)
            }
        }()

        // create a new waitgroup (count: 0)
        var wg sync.WaitGroup

        // call done creating a negative
        // waitgroup counter
        wg.Done()

        // this line is never reached
        fmt.Println("waitgroup is unblocked")
    }()
}
```

```
$ go test -v -timeout 1s

=== RUN   Test_WaitGroup_Done
=== PAUSE Test_WaitGroup_Done
=== CONT Test_WaitGroup_Done
done_test.go:20: sync: negative WaitGroup counter
--- FAIL: Test_WaitGroup_Done (0.00s)
FAIL
exit status 1
FAIL    demo    0.416s
```

```
Go Version: go1.19
```

Wrapping Up Wait Groups

Using a `sync.WaitGroup` is a great way to manage the number of goroutines or any other number of tests that need to finish before your program can continue.

As you can see, we can effectively use a `sync.WaitGroup` to manage the thumbnail generator goroutines from our initial example.

In Listing 13.13, we create a new `sync.WaitGroup`. Then, in the `for` loop, we use the `sync.WaitGroup.Add` method to add 1 to the `sync.WaitGroup`. We then pass a pointer to the `generateThumbnail` function to `sync.WaitGroup`. A pointer is needed because the `generateThumbnail` function needs to be able to modify the `sync.WaitGroup` by calling the `sync.WaitGroup.Done` method.

Finally, we call the `sync.WaitGroup.Wait` method to wait for all of the goroutines to finish.

Listing 13.13 Using a `sync.WaitGroup` to Manage the Thumbnail Generator Goroutines

```
func Test_ThumbnailGenerator(t *testing.T) {
    t.Parallel()

    // image that we need thumbnails for
    const image = "foo.png"

    var wg sync.WaitGroup

    // start 5 goroutines to generate thumbnails
    for i := 0; i < 5; i++ {
        wg.Add(1)

        // start a new goroutine for each thumbnail
        go generateThumbnail(&wg, image, i+1)
    }

    fmt.Println("Waiting for thumbnails to be generated")

    // wait for all goroutines to finish
    wg.Wait()

    fmt.Println("Finished generate all thumbnails")
}
```

The `generateThumbnail` function now receives a pointer to the `sync.WaitGroup` and defers a call to the `sync.WaitGroup.Done` method to indicate that the goroutine is finished when the function exits.

Finally, as you can see from our test output in Listing 13.14, the application now finishes successfully.

Listing 13.14 Generating Thumbnails Using a `sync.WaitGroup`

```
func generateThumbnail(wg *sync.WaitGroup, image string, size int) {
    defer wg.Done()

    // thumbnail to be generated
    thumb := fmt.Sprintf("%s@%dx.png", image, size)

    fmt.Println("Generating thumbnail:", thumb)

    // wait for the thumbnail to be ready
    time.Sleep(time.Millisecond * time.Duration(size))

    fmt.Println("Finished generating thumbnail:", thumb)
}
```

```
$ go test -v

=== RUN   Test_ThumbnailGenerator
=== PAUSE Test_ThumbnailGenerator
=== CONT Test_ThumbnailGenerator
Waiting for thumbnails to be generated
Generating thumbnail: foo.png@5x.png
Generating thumbnail: foo.png@3x.png
Generating thumbnail: foo.png@4x.png
Generating thumbnail: foo.png@2x.png
Generating thumbnail: foo.png@1x.png
Finished generating thumbnail: foo.png@1x.png
Finished generating thumbnail: foo.png@2x.png
Finished generating thumbnail: foo.png@3x.png
Finished generating thumbnail: foo.png@4x.png
Finished generating thumbnail: foo.png@5x.png
Finished generate all thumbnails
--- PASS: Test_ThumbnailGenerator (0.01s)
PASS
ok      demo    0.310s
```

```
Go Version: go1.19
```

Error Management with Error Groups

One of the downsides of the `sync.WaitGroup` is that it has no error management built in to capture errors that occur in the goroutines. It also has an API that requires exact implementation; otherwise, it panics.

To help address some of these issues, the [golang.org/x/sync/errgroup](https://pkg.go.dev/golang.org/x/sync/errgroup),¹⁰ Listing 13.15, package was introduced, providing a simpler API as well as built-in error management.

Listing 13.15 The [golang.org/x/sync/errgroup](https://pkg.go.dev/golang.org/x/sync/errgroup) Package

```
$ go doc golang.org/x/sync/errgroup

package errgroup // import "golang.org/x/sync/errgroup"

Package errgroup provides synchronization, error propagation, and Context
▶cancellation for groups of goroutines working on subtasks of a common task.

type Group struct{ ... }
    func WithContext(ctx context.Context) (*Group, context.Context)

```

Go Version: go1.19

The Problem

Consider Listing 13.16. In the example, a number of goroutines are launched to call the `generateThumbnail` functions. A `sync.WaitGroup` is used to wait for all the goroutines to finish.

Listing 13.16 Managing Goroutines with `sync.WaitGroup`

```
func Test_ThumbnailGenerator(t *testing.T) {
    t.Parallel()

    // image that we need thumbnails for
    const image = "foo.png"

    var wg sync.WaitGroup

    // start 5 goroutines to generate thumbnails
    for i := 0; i < 5; i++ {
        wg.Add(1)

        // start a new goroutine for each thumbnail
        go generateThumbnail(&wg, image, i+1)
    }
}
```

10. <https://pkg.go.dev/golang.org/x/sync/errgroup>

```

    fmt.Println("Waiting for thumbnails to be generated")

    // wait for all goroutines to finish
    wg.Wait()

    fmt.Println("Finished generate all thumbnails")
}

```

Inside the `generateThumbnail` function, Listing 13.17, we see that there is an error that occurs if the `size` argument is divisible by 5, which causes a panic, and the application crashes.

Listing 13.17 The generateThumbnail Function

```

func generateThumbnail(wg *sync.WaitGroup, image string, size int) {
    defer wg.Done()

    // error if the size is divisible by 5
    if size%5 == 0 {
        // how do we return this error back to the main
        // goroutine without panicking?
        err := fmt.Errorf("%d is divisible by 5", size)
        panic(err)
    }

    // thumbnail to be generated
    thumb := fmt.Sprintf("%s@%dx.png", image, size)

    fmt.Println("Generating thumbnail:", thumb)

    // wait for the thumbnail to be ready
    time.Sleep(time.Millisecond * time.Duration(size))

    fmt.Println("Finished generating thumbnail:", thumb)
}

```

```

$ go test -v

=== RUN   Test_ThumbnailGenerator
=== PAUSE Test_ThumbnailGenerator
=== CONT Test_ThumbnailGenerator
Waiting for thumbnails to be generated
panic: 5 is divisible by 5

```

(continued)

```

goroutine 9 [running]:
demo.generateThumbnail(0x0?, {0x10268ac77?, 0x0?}, 0x0?)
    ./demo_test.go:47 +0x21c
created by demo.Test_ThumbnailGenerator
    ./demo_test.go:24 +0x4c
exit status 2
FAIL    demo    0.619s

```

```
Go Version: go1.19
```

The Error Group

The `errgroup.Group`¹¹ type provides a minimal API when compared to the `sync.WaitGroup` type. There are only two methods: `errgroup.Group.Go`¹² and `errgroup.Group.Wait`.¹³

The `errgroup.Group` type, Listing 13.18, manages the counter for you, so there is no need for counter management as there is with the `sync.WaitGroup` type.

Listing 13.18 The `errgroup.Group` Type

```

$ go doc golang.org/x/sync/errgroup.Group

package errgroup // import "golang.org/x/sync/errgroup"

type Group struct {
    // Has unexported fields.
}

A Group is a collection of goroutines working on subtasks that are part
➔of the same overall task.

A zero Group is valid and does not cancel on error.

func WithContext(ctx context.Context) (*Group, context.Context)
func (g *Group) Go(f func() error)
func (g *Group) Wait() error

```

```
Go Version: go1.19
```

11. <https://pkg.go.dev/golang.org/x/sync/errgroup#Group>

12. <https://pkg.go.dev/golang.org/x/sync/errgroup#Group.Go>

13. <https://pkg.go.dev/golang.org/x/sync/errgroup#Group.Wait>

The Go Method

You use the `errgroup.Group.Go` method, Listing 13.19, to launch a goroutine for the provided `func() error` function provided. The `func() error` function is executed in a goroutine. If the function returns an error, the `errgroup.Group` type captures the error, cancels the other goroutines, and returns the error to the caller from the `errgroup.Group.Wait` method.

Listing 13.19 The `errgroup.Group.Go` Method

```
$ go doc golang.org/x/sync/errgroup.Group.Go

package errgroup // import "golang.org/x/sync/errgroup"

func (g *Group) Go(f func() error)
    Go calls the given function in a new goroutine.

    The first call to return a non-nil error cancels the group; its error
    will be returned by Wait.
```

```
Go Version: go1.19
```

The Wait Method

You use `errgroup.Group.Wait` method, Listing 13.20, to wait for all the goroutines to finish. If any of the goroutines return an error, the `errgroup.Group` type returns the error to the caller.

Listing 13.20 The `errgroup.Group.Wait` Method

```
$ go doc golang.org/x/sync/errgroup.Group.Wait

package errgroup // import "golang.org/x/sync/errgroup"

func (g *Group) Wait() error
    Wait blocks until all function calls from the Go method have returned,
    then returns the first non-nil error (if any) from them.
```

```
Go Version: go1.19
```

It is important to understand the `errgroup.Group.Wait` method returns *only* the *first* error that occurs. Any other errors that occur are ignored. In Listing 13.21, we are calling the `Go` method on an `errgroup.Group` 10 times. Each time we pass a function that sleeps for a random time, prints a message, and then returns an error. From the output of the test, you can see that the error returned by the `Wait` method was returned from function number 4. The other nine errors are discarded.

Listing 13.21 The `errgroup.Group.Wait` Method Returns Only the First Error

```

func Test_ErrorGroup_Multiple_Errors(t *testing.T) {
    t.Parallel()

    var wg errgroup.Group

    for i := 0; i < 10; i++ {
        i := i + 1

        wg.Go(func() error {

            time.Sleep(time.Millisecond * time.Duration(rand.Intn(10)))

            fmt.Printf("about to error from %d\n", i)

            return fmt.Errorf("error %d", i)

        })
    }

    err := wg.Wait()
    if err != nil {
        t.Fatal(err)
    }
}

```

```

$ go test -v

=== RUN   Test_ErrorGroup_Multiple_Errors
=== PAUSE Test_ErrorGroup_Multiple_Errors
=== CONT  Test_ErrorGroup_Multiple_Errors
about to error from 4
about to error from 6
about to error from 3
about to error from 1
about to error from 9
about to error from 10
about to error from 8
about to error from 5
about to error from 2
about to error from 7
demo_test.go:38: error 4

```

```

--- FAIL: Test_ErrorGroup_Multiple_Errors (0.01s)
FAIL
exit status 1
FAIL    demo    0.263s

```

```

Go Version: go1.19

```

Listening for Error Group Cancellation

When launching a number of goroutines, it can often be useful to let others know the tasks have all completed. The `errgroup.WithContext`¹⁴ function, Listing 13.22, returns a new `errgroup.Group`, as well as a `context.Context` that can be listened to for cancellation.

Listing 13.22 The `errgroup.WithContext` Function

```

func Test_ErrorGroup_Context(t *testing.T) {
    t.Parallel()

    // create a new error group
    // and a context that will be canceled
    // when the group is done
    wg, ctx := errgroup.WithContext(context.Background())

    // create a quit channel for the goroutine
    // waiting for the context to be canceled
    // can close to signal the goroutine has finished
    quit := make(chan struct{})

    // launch a goroutine that will
    // wait for the errgroup context to finish
    go func() {
        fmt.Println("waiting for context to cancel")

        // wait for the context to be canceled
        <-ctx.Done()

        fmt.Println("context canceled")

        // close the quit channel so the test
        // will finish
        close(quit)
    }()
}

```

(continued)

14. <https://pkg.go.dev/golang.org/x/sync/errgroup#WithContext>

```

    }()

    // add a task to the errgroup
    wg.Go(func() error {
        time.Sleep(time.Millisecond * 5)
        return nil
    })

    // wait for the errgroup to finish
    err := wg.Wait()
    if err != nil {
        t.Fatal(err)
    }

    // wait for the context goroutine to finish
    <-quit
}

```

```

$ go test -v

=== RUN   Test_ErrorGroup_Context
=== PAUSE Test_ErrorGroup_Context
=== CONT Test_ErrorGroup_Context
waiting for context to cancel
context canceled
--- PASS: Test_ErrorGroup_Context (0.01s)
PASS
ok      demo    0.725s

```

```
Go Version: go1.19
```

Wrapping Up Error Groups

The `errgroup.Group` type provides a simpler API than the `sync.WaitGroup` type. It also provides built-in error management. This makes managing goroutines with error handling much easier. The downside is that it is not as flexible as the `sync.WaitGroup` type, where you are in charge of managing the counter.

Which type you choose to use will vary from situation to situation, so it is important to understand the tradeoffs and benefits of each before deciding which type to use.

Using a `errgroup.Group` allows you to clean up your code significantly to make it easier to understand and to manage errors.

As you can see from Listing 13.23, the `generateThumbnail` function no longer needs to take a `sync.WaitGroup` as an argument.

Listing 13.23 Update the `generateThumbnail` Function to use `errgroup.Group`

```
func generateThumbnail(image string, size int) error {  
  
    // error if the size is divisible by 5  
    if size%5 == 0 {  
        return fmt.Errorf("%d is divisible by 5", size)  
    }  
  
    // thumbnail to be generated  
    thumb := fmt.Sprintf("%s@%dx.png", image, size)  
  
    fmt.Println("Generating thumbnail:", thumb)  
  
    // wait for the thumbnail to be ready  
    time.Sleep(time.Millisecond * time.Duration(size))  
  
    fmt.Println("Finished generating thumbnail:", thumb)  
  
    return nil  
}
```

Being able to return an error from the function means the function no longer needs to panic.

As you can see from the output, Listing 13.24, the `generateThumbnail` function no longer panics, and the test is now able to exit properly.

Listing 13.24 Using the `errgroup.Group` Type

```
func Test_ThumbnailGenerator(t *testing.T) {  
    t.Parallel()  
  
    // image that we need thumbnails for  
    const image = "foo.png"  
  
    // create a new error group  
    var wg errgroup.Group  
  
    // start 5 goroutines to generate thumbnails  
    for i := 0; i < 5; i++ {  
  
        // capture the i to the current scope  
        i := i
```

(continued)

```

    // start a new goroutine for each thumbnail
    wg.Go(func() error {

        // return the result of generateThumbnail
        return generateThumbnail(image, i)
    })

}

fmt.Println("Waiting for thumbnails to be generated")

// wait for all goroutines to finish
err := wg.Wait()

// check for any errors
if err != nil {
    t.Fatal(err)
}

fmt.Println("Finished generate all thumbnails")
}

```

```

$ go test -v

=== RUN   Test_ThumbnailGenerator
=== PAUSE Test_ThumbnailGenerator
=== CONT Test_ThumbnailGenerator
Waiting for thumbnails to be generated
Generating thumbnail: foo.png@4x.png
Generating thumbnail: foo.png@3x.png
Generating thumbnail: foo.png@1x.png
Generating thumbnail: foo.png@2x.png
Finished generating thumbnail: foo.png@1x.png
Finished generating thumbnail: foo.png@2x.png
Finished generating thumbnail: foo.png@3x.png
Finished generating thumbnail: foo.png@4x.png
demo_test.go:43: 0 is divisible by 5
--- FAIL: Test_ThumbnailGenerator (0.00s)
FAIL
exit status 1
FAIL    demo    0.570s

```

```
Go Version: go1.19
```

Data Races

When writing concurrent applications, it is common to run into what is called a race condition.¹⁵ A race condition occurs when two different goroutines try to access the same shared resource.

Consider Listing 13.25. There are two different goroutines. One goroutine inserts values into a map, and the other goroutine ranges over the map and prints the values.

Listing 13.25 Two Goroutines Accessing a Shared Map

```
// launch a goroutine to
// write data in the map
go func() {
    for i := 0; i < 10; i++ {

        // loop putting data in the map
        data[i] = true
    }

    // cancel the context
    cancel()
}()

// launch a goroutine to
// read data from the map
go func() {
    // loop through the map
    // and print the keys/values
    for k, v := range data {
        fmt.Printf("%d: %v\n", k, v)
    }
}()
```

In Listing 13.26, we use those two goroutines to write a test to assert the map is written to and read from correctly.

Listing 13.26 A Passing Testing Without the Race Detector

```
func Test_Mutex(t *testing.T) {
    t.Parallel()
```

(continued)

15. https://en.wikipedia.org/wiki/Race_condition

```

// create a new cancellable context
// to stop the test when the goroutines
// are finished
ctx := context.Background()
ctx, cancel := context.WithTimeout(ctx, 20*time.Millisecond)
defer cancel()

// create a map to be used
// as a shared resource
data := map[int]bool{}

// launch a goroutine to
// write data in the map
go func() {
    for i := 0; i < 10; i++ {

        // loop putting data in the map
        data[i] = true
    }

    // cancel the context
    cancel()
}()

// launch a goroutine to
// read data from the map
go func() {
    // loop through the map
    // and print the keys/values
    for k, v := range data {
        fmt.Printf("%d: %v\n", k, v)
    }
}()

// wait for the context to be canceled
<-ctx.Done()

if len(data) != 10 {
    t.Fatalf("expected 10 items in the map, got %d", len(data))
}
}

```

```

$ go test -v

=== RUN   Test_Mutex
=== PAUSE Test_Mutex
=== CONT  Test_Mutex
--- PASS: Test_Mutex (0.00s)
PASS
ok      demo    0.471s

```

```

Go Version: go1.19

```

A quick glance at the test output, Listing 13.26, would seem to imply that the tests have passed successfully, but this is not the case.

The Race Detector

A few of the Go commands, such as `test` and `build`, have the `-race` flag exposed. When used, the `-race` flag tells the Go compiler to create a special version of the binary or test binary that will detect and report race conditions.

If we run the test again, this time with the `-race` flag, we get a *very* different result, as shown in Listing 13.27.

Listing 13.27 Tests Failing with the Race Detector

```

$ go test -v -race

=== RUN   Test_Mutex
=== PAUSE Test_Mutex
=== CONT  Test_Mutex
--- PASS: Test_Mutex (0.00s)
=====
WARNING: DATA RACE
Read at 0x00c00011c3f0 by goroutine 9:
  runtime.mapdelete()
      /usr/local/go/src/runtime/map.go:695 +0x46c
demo.Test_Mutex.func2()
  ./demo_test.go:46 +0x50

Previous write at 0x00c00011c3f0 by goroutine 8:
  runtime.mapaccess2_fast64()
      /usr/local/go/src/runtime/map_fast64.go:53 +0x1cc
demo.Test_Mutex.func1()
  ./demo_test.go:32 +0x50

```

(continued)


```

Goroutine 9 (running) created at:
  demo.Test_Mutex()
    ./demo_test.go:43 +0x188
  testing.tRunner()
    /usr/local/go/src/testing/testing.go:1439 +0x18c
  testing.(*T).Run.func1()
    /usr/local/go/src/testing/testing.go:1486 +0x44

Goroutine 8 (finished) created at:
  demo.Test_Mutex()
    ./demo_test.go:28 +0x124
  testing.tRunner()
    /usr/local/go/src/testing/testing.go:1439 +0x18c
  testing.(*T).Run.func1()
    /usr/local/go/src/testing/testing.go:1486 +0x44
=====
FAIL
exit status 1
FAIL    demo    0.962s

```

```

Go Version: go1.19

```

As you can see from the output, the Go race detector found a race condition in our code.

If we examine the top two entries in a race condition warning, Listing 13.28, it tells us where the two conflicting lines of code are.

Listing 13.28 Reading the Race Detector Output

```

Read at 0x00c00018204b by goroutine 9:
  demo.Test_Mutex.func2()
    problem/demo_test.go:46 +0xa5

Previous write at 0x00c00018204b by goroutine 8:
  demo.Test_Mutex.func1()
    problem/demo_test.go:32 +0x5c

```

A read of the shared resource was happening at `demo_test.go:46`, and a write was happening at `demo_test.go:32`. We need to synchronize or lock these two goroutines so that they don't both try to access the shared resource at the same time.

Most, but Not All

The Go race detector makes a simple guarantee with you (the end user).

The Go race detector may not find *all* the race conditions in your code, but the ones it does find are *real* and *must* be fixed.

A race condition *will* panic and crash your application. If the race detector finds a race condition, you *must* fix it.

Wrapping Up the Race Detector

The race detector is an *invaluable* tool when developing Go applications. When running tests with the `-race` flag, you will notice a slowdown in test performance. The race detector has to do a lot of work to track those conditions.

Always enable the `-race` flag on your CI, such as GitHub Actions.

Once identified, the `sync` package, Listing 13.29, provides a number of ways that you can fix issues.

Listing 13.29 The `sync` Package

```
$ go doc -short sync

type Cond struct{ ... }
    func NewCond(1 Locker) *Cond
type Locker interface{ ... }
type Map struct{ ... }
type Mutex struct{ ... }
type Once struct{ ... }
type Pool struct{ ... }
type RWMutex struct{ ... }
type WaitGroup struct{ ... }
```

```
Go Version: go1.19
```

Synchronizing Access with a Mutex

When you run tests with the `-race` flag, Go's built-in race detector can help you find data races in your code. In Listing 13.30, for example, tests that pass normally fail when run with the race detector. The failure message lists the data race found and where the reads and writes occur in our code.

Listing 13.30 Detecting Race Conditions in Tests

```
$ go test -v -race

=== RUN   Test_Mutex
=== PAUSE Test_Mutex
```

(continued)

```

=== CONT Test_Mutex
--- PASS: Test_Mutex (0.00s)
PASS
=====
WARNING: DATA RACE
Read at 0x00c00019e3f0 by goroutine 9:
  runtime.mapdelete()
    /usr/local/go/src/runtime/map.go:695 +0x46c
demo.Test_Mutex.func2()
  ./demo_test.go:46 +0x50

Previous write at 0x00c00019e3f0 by goroutine 8:
  runtime.mapaccess2_fast64()
    /usr/local/go/src/runtime/map_fast64.go:53 +0x1cc
demo.Test_Mutex.func1()
  ./demo_test.go:32 +0x50

Goroutine 9 (running) created at:
  demo.Test_Mutex()
    ./demo_test.go:43 +0x188
  testing.tRunner()
    /usr/local/go/src/testing/testing.go:1439 +0x18c
  testing.(*T).Run.func1()
    /usr/local/go/src/testing/testing.go:1486 +0x44

Goroutine 8 (finished) created at:
  demo.Test_Mutex()
    ./demo_test.go:28 +0x124
  testing.tRunner()
    /usr/local/go/src/testing/testing.go:1439 +0x18c
  testing.(*T).Run.func1()
    /usr/local/go/src/testing/testing.go:1486 +0x44
=====
0: true
2: true
4: true
7: true
9: true
1: true
3: true
5: true
6: true
8: true

```

```
Found 1 data race(s)
exit status 66
FAIL    demo    0.862s
```

```
Go Version: go1.19
```

In the test in Listing 13.31, we have two different goroutines. The first is modifying a shared resource—in this case, a map. The second goroutine is ranging over the map and printing out the map's values.

In order for us to be able to fix this race condition, we need to be able to synchronize access to the shared resource.

Listing 13.31 Two Goroutines Accessing a Shared Map

```
// launch a goroutine to
// read data from the map
go func() {
    // loop through the map
    // and print the keys/values
    for k, v := range data {
        fmt.Printf("%d: %v\n", k, v)
    }
}()

// launch a goroutine to
// put data in the map
go func() {
    for i := 0; i < 10; i++ {

        // loop putting data in the map
        data[i] = true
    }

    // cancel the context
    cancel()
}()
```

Locker

To synchronize access to the shared resource, we need to be able to **lock** access to the resource. By locking a shared resource, we can ensure that only one goroutine at a time can access the resource and that the resource is not modified by another goroutine while it is locked.

The `sync.Locker`¹⁶ interface, Listing 13.32, defines the methods that a type must implement to be able to lock and unlock a shared resource.

Listing 13.32 The `sync.Locker` Interface

```
$ go doc sync.Locker

package sync // import "sync"

type Locker interface {
    Lock()
    Unlock()
}

A Locker represents an object that can be locked and unlocked.
```

Go Version: go1.19

Locker Methods

You can use the `sync.Locker.Lock`¹⁷ method, Listing 13.33, to lock the shared resource. Once a resource is locked, no other goroutine can access the resource until it is unlocked.

Listing 13.33 The `sync.Locker.Lock` Method

```
$ go doc sync.Mutex.Lock

package sync // import "sync"

func (m *Mutex) Lock()
    Lock locks m. If the lock is already in use, the calling goroutine blocks
    →until the mutex is available.
```

Go Version: go1.19

You can use the `sync.Locker.Unlock`¹⁸ method, Listing 13.34, to unlock the shared resource. Once a resource is unlocked, other goroutines can access the resource.

16. <https://pkg.go.dev/sync#Locker>

17. <https://pkg.go.dev/sync#Locker.Lock>

18. <https://pkg.go.dev/sync#Locker.Unlock>

Listing 13.34 The `sync.Locker.Unlock` Method

```
$ go doc sync.Mutex.Unlock

package sync // import "sync"

func (m *Mutex) Unlock()
    Unlock unlocks m. It is a run-time error if m is not locked on entry to
    ➤Unlock.

    A locked Mutex is not associated with a particular goroutine. It is
    ➤allowed for one goroutine to lock a Mutex and then arrange for another
    ➤goroutine to unlock it.
```

Go Version: go1.19

Using a Mutex

The most basic mutex available in Go is the `sync.Mutex` type, Listing 13.35. The `sync.Mutex` uses a basic binary semaphore lock. This means that only one goroutine can access the resource at a time.

Listing 13.35 The `sync.Mutex` Type

```
$ go doc sync.Mutex

package sync // import "sync"

type Mutex struct {
    // Has unexported fields.
}
    ➤A Mutex is a mutual exclusion lock. The zero value for a Mutex is an
    ➤unlocked mutex.

    A Mutex must not be copied after first use.

func (m *Mutex) Lock()
func (m *Mutex) TryLock() bool
func (m *Mutex) Unlock()
```

Go Version: go1.19

To use a `sync.Mutex`, you need to wrap the areas of code that you want to synchronize access to by first locking the `sync.Mutex` and then unlocking it. For example, in the

second goroutine in Listing 13.36, a mutex is being used to lock access around writing values into the `data` map.

Listing 13.36 Locking Resources with a `sync.Mutex`

```
// launch a goroutine to
// read data from the map
go func() {
    // lock the mutex
    mu.Lock()

    // loop through the map
    // and print the keys/values
    for k, v := range data {
        fmt.Printf("%d: %v\n", k, v)
    }

    // unlock the mutex
    mu.Unlock()
}()

// launch a goroutine to
// put data in the map
go func() {
    for i := 0; i < 10; i++ {
        // lock the mutex
        mu.Lock()

        // loop putting data in the map
        data[i] = true

        // unlock the mutex
        mu.Unlock()
    }

    // cancel the context
    cancel()
}()
```

By locking access to the shared resource, you can ensure that only one goroutine at a time can access the resource. Our test output, Listing 13.37, confirms that the shared resource is only accessed by one goroutine at a time with a successful exit.

Listing 13.37 Passing Race Detector Tests

```
$ go test -v -race

=== RUN   Test_Mutex
=== PAUSE Test_Mutex
=== CONT  Test_Mutex
9: true
--- PASS: Test_Mutex (0.00s)
0: true
2: true
5: true
8: true
7: true
1: true
PASS
3: true
4: true
6: true
ok      demo      0.810s
```

```
Go Version: go1.19
```

RWMutex

Often, applications read a shared resource instead of writing to them. The `sync.Mutex` is a very heavy-weight locking mechanism. Access to a shared resource, whether it be a read or write, is blocked until the resource is unlocked. Only one goroutine can access a shared resource at a time.

When you want to be able to both read and write to a shared resource, you need to use a `sync.RWMutex`, Listing 13.38. The `sync.RWMutex` is a lighter-weight locking mechanism. A `sync.RWMutex` can allow many goroutines to read from the resource at the same time, but only one goroutine can write to the resource at a time.

Listing 13.38 The `sync.RWMutex` Type

```
$ go doc sync.RWMutex

package sync // import "sync"

type RWMutex struct {
    // Has unexported fields.
}
```

(continued)

A `RWMutex` is a reader/writer mutual exclusion lock. The lock can be held
 ↳by an arbitrary number of readers or a single writer. The zero value for
 ↳a `RWMutex` is an unlocked mutex.

A `RWMutex` must not be copied after first use.

If a goroutine holds a `RWMutex` for reading and another goroutine might
 ↳call `Lock`, no goroutine should expect to be able to acquire a read lock
 ↳until the initial read lock is released. In particular, this prohibits
 ↳recursive read locking. This is to ensure that the lock eventually
 ↳becomes available; a blocked `Lock` call excludes new readers from
 ↳acquiring the lock.

```
func (rw *RWMutex) Lock()
func (rw *RWMutex) RLock()
func (rw *RWMutex) RLocker() Locker
func (rw *RWMutex) RUnlock()
func (rw *RWMutex) TryLock() bool
func (rw *RWMutex) TryRLock() bool
func (rw *RWMutex) Unlock()
```

Go Version: go1.19

The `sync.RWMutex` offers two additional methods beyond those of the `sync.Locker` interface. You can use the `sync.RWMutex.RLock`¹⁹ and `sync.RWMutex.RUnlock`²⁰ methods to lock the resource for reading. The `sync.RWMutex.Lock`²¹ and `sync.RWMutex.Unlock`²² methods are used to lock the resource across all goroutines for writing.

In Listing 13.39, we update the goroutine that is reading the resource to use the `sync.RWMutex.RLock` method instead of the `sync.Mutex.Lock` method. This will allow for many goroutines to read from the resource at the same time.

Listing 13.39 Using a `sync.RWMutex`

```
// launch a goroutine to
// read data from the map
go func() {
    // lock the mutex
    mu.RLock()
```

19. <https://pkg.go.dev/sync#RWMutex.RLock>

20. <https://pkg.go.dev/sync#RWMutex.RUnlock>

21. <https://pkg.go.dev/sync#RWMutex.Lock>

22. <https://pkg.go.dev/sync#RWMutex.Unlock>

```
// loop through the map
// and print the keys/values
for k, v := range data {
    fmt.Printf("%d: %v\n", k, v)
}

// unlock the mutex
mu.RUnlock()
}()
```

```
$ go test -v -race

=== RUN   Test_RWMutex
=== PAUSE Test_RWMutex
=== CONT Test_RWMutex
4: true
5: true
6: true
--- PASS: Test_RWMutex (0.00s)
7: true
9: true
0: true
1: true
PASS
8: true
2: true
3: true
ok      demo    0.917s
```

```
Go Version: go1.19
```

The tests in Listing 13.39 continue to pass, but the performance of the program is improved by allowing multiple goroutines to read from the shared resource instead of arbitrarily locking all goroutines all at once.

Improper Usage

When using either `sync.Mutex` or `sync.RWMutex`, you must take care in making sure to lock and unlock in the proper order.

Consider Listing 13.40. We have a `sync.Mutex`, and we attempt to call `sync.Mutex.Lock` twice.

Listing 13.40 Attempting to Lock a `sync.Mutex` Twice

```

func Test_Mutex_Locks(t *testing.T) {
    t.Parallel()

    // create a new mutex
    var mu sync.Mutex

    // lock the mutex
    mu.Lock()

    fmt.Println("locked. locking again.")

    // try to lock the mutex again
    // this will block/deadlock
    // because the mutex is already locked
    // and the lock was not released
    mu.Lock()

    fmt.Println("unlocked twice")
}

```

The result is the program will deadlock and crash, as shown in Listing 13.41. The reason is that a call to `sync.Mutex.Lock` blocks until the `sync.Mutex.Unlock` method is called. Because we have already locked the `sync.Mutex`, the second call to `sync.Mutex.Lock` blocks indefinitely because it is never unlocked.

Listing 13.41 A Panic while Trying to Unlock an Already-Unlocked `sync.Mutex`

```

$ go test -v -timeout 10ms

=== RUN   Test_Mutex_Locks
=== PAUSE Test_Mutex_Locks
=== CONT Test_Mutex_Locks
locked. locking again.
panic: test timed out after 10ms

goroutine 33 [running]:
testing.(*M).startAlarm.func1()
    /usr/local/go/src/testing/testing.go:2029 +0x8c
created by time.goFunc
    /usr/local/go/src/time/sleep.go:176 +0x3c

```

```

goroutine 1 [chan receive]:
testing.tRunner.func1()
    /usr/local/go/src/testing/testing.go:1405 +0x45c
testing.tRunner(0x140001361a0, 0x1400010fcb8)
    /usr/local/go/src/testing/testing.go:1445 +0x14c
testing.runTests(0x1400001e1e0?, {0x100ec9ea0, 0x1, 0x1},
    ↪{0xe000000000000000?, 0x100cf5218?, 0x100ed2640?})
    /usr/local/go/src/testing/testing.go:1837 +0x3f0

testing.(*M).Run(0x1400001e1e0)
    /usr/local/go/src/testing/testing.go:1719 +0x500
main.main()
    _testmain.go:47 +0x1d0

goroutine 4 [semacquire]:
sync.runtime_SemacquireMutex(0x1400000e018?, 0x20?, 0x17?)
    /usr/local/go/src/runtime/sema.go:71 +0x28
sync.(*Mutex).lockSlow(0x14000012140)
    /usr/local/go/src/sync/mutex.go:162 +0x180
sync.(*Mutex).Lock(...)
    /usr/local/go/src/sync/mutex.go:81
demo.Test_Mutex_Locks(0x0?)
    ./demo_test.go:25 +0x130
testing.tRunner(0x14000136340, 0x100e25298)
    /usr/local/go/src/testing/testing.go:1439 +0x110
created by testing.(*T).Run
    /usr/local/go/src/testing/testing.go:1486 +0x300
exit status 2
FAIL    demo    0.527s

```

Go Version: go1.19

Worse than a deadlock caused by waiting for a lock that will never be unlocked is unlocking a lock that has not been locked.

In Listing 13.42, the result is a fatal error that crashes the application. The reason is that we have not locked the `sync.Mutex` before attempting to unlock it.

Listing 13.42 A Panic while Trying to Unlock an Unlocked `sync.Mutex`

```

func Test_Mutex_Unlock(t *testing.T) {
    t.Parallel()

    // create a new mutex
    var mu sync.Mutex

```

(continued)

```

    // unlock the mutex
    mu.Unlock()
}

```

```

$ go test -v

=== RUN   Test_Mutex_Unlock
=== PAUSE Test_Mutex_Unlock
=== CONT Test_Mutex_Unlock
fatal error: sync: unlock of unlocked mutex

goroutine 18 [running]:
runtime.throw(0x104573b02?, 0x1400005af18?)
    /usr/local/go/src/runtime/panic.go:992 +0x50 fp=0x1400005aee0
    ↪sp=0x1400005aeb0 pc=0x1044ba9a0
sync.throw(0x104573b02?, 0x1045b6260?)
    /usr/local/go/src/runtime/panic.go:978 +0x24 fp=0x1400005af00
    ↪sp=0x1400005aee0 pc=0x1044e5664
sync.(*Mutex).unlockSlow(0x140001280b0, 0xffffffff)
    /usr/local/go/src/sync/mutex.go:220 +0x3c fp=0x1400005af30
    ↪sp=0x1400005af00 pc=0x1044ef44c
sync.(*Mutex).Unlock(...)
    /usr/local/go/src/sync/mutex.go:214
demo.Test_Mutex_Unlock(0x0?)
    ./demo_test.go:16 +0x74 fp=0x1400005af60 sp=0x1400005af30
    ↪pc=0x10456d794
testing.tRunner(0x1400010b380, 0x1045c9298)
    /usr/local/go/src/testing/testing.go:1439 +0x110 fp=0x1400005afb0
    ↪sp=0x1400005af60 pc=0x104537660
testing.(*T).Run.func1()
    /usr/local/go/src/testing/testing.go:1486 +0x30 fp=0x1400005afd0
    ↪sp=0x1400005afb0 pc=0x1045383d0
runtime.goexit()
    /usr/local/go/src/runtime/asm_arm64.s:1263 +0x4 fp=0x1400005afd0
    ↪sp=0x1400005afd0 pc=0x1044ea2a4
created by testing.(*T).Run
    /usr/local/go/src/testing/testing.go:1486 +0x300

goroutine 1 [chan receive]:
testing.tRunner.func1()
    /usr/local/go/src/testing/testing.go:1405 +0x45c
testing.tRunner(0x1400010b1e0, 0x14000131cb8)
    /usr/local/go/src/testing/testing.go:1445 +0x14c
testing.runTests(0x140001421e0?, {0x10466dea0, 0x1, 0x1}, {0xa500000000000000?,

```

```

↳0x104499218?, 0x104676640?})
    /usr/local/go/src/testing/testing.go:1837 +0x3f0
testing.(*M).Run(0x140001421e0)
    /usr/local/go/src/testing/testing.go:1719 +0x500
main.main()
    _testmain.go:47 +0x1d0

exit status 2
FAIL    demo    0.260s

```

```

Go Version: go1.19

```

Wrapping Up Read/Write Mutexes

While there are pitfalls and areas of concern when using mutexes, such as deadlocks and improper usage, `sync.Mutex` and `sync.RWMutex` are excellent tools for protecting shared resources. They are also the most commonly used locking mechanisms in Go.

Performing Tasks Only Once

There are many times when you want to perform a task only once. For example, you might want to create a database connection only once and then use it to perform a number of queries. You can use the `sync.Once`²³ type to do this.

As you can see from the documentation in Listing 13.43, the use of `sync.Once` is very simple. You just need to create a variable of type `sync.Once` and then call the `sync.Once.Do`²⁴ method with a function that you want to run only once.

Listing 13.43 The `sync.Once` Type

```

$ go doc -all sync.Once

package sync // import "sync"

type Once struct {
    // Has unexported fields.
}

Once is an object that will perform exactly one action.

A Once must not be copied after first use.

```

(continued)

23. <https://pkg.go.dev/sync#Once>

24. <https://pkg.go.dev/sync#Once.Do>

```
func (o *Once) Do(f func())
    Do calls the function f if and only if Do is being called for the first
    ↪time for this instance of Once. In other words, given

        var once Once

    if once.Do(f) is called multiple times, only the first call will invoke f,
    ↪even if f has a different value in each invocation. A new instance of
    ↪Once is required for each function to execute.

    Do is intended for initialization that must be run exactly once. Since f
    ↪is niladic, it may be necessary to use a function literal to capture the
    ↪arguments to a function to be invoked by Do:

        config.once.Do(func() { config.init(filename) })

    Because no call to Do returns until the one call to f returns, if f
    ↪causes Do to be called, it will deadlock.

    If f panics, Do considers it to have returned; future calls of Do return
    ↪without calling f.
```

Go Version: go1.19

The Problem

Often we want to use `sync.Once` to perform some heavy, expensive tasks only once.

Consider Listing 13.44. The `Build` function can be called many times, but we only want it to run once because it takes some time to complete.

Listing 13.44 The `Build` Method Is Slow and Should Be Called Only Once

```
type Builder struct {
    Built bool
}

func (b *Builder) Build() error {

    fmt.Print("building...")

    time.Sleep(10 * time.Millisecond)

    fmt.Println("built")
```

```

    b.Built = true

    // validate the message
    if !b.Built {
        return fmt.Errorf("expected builder to be built")
    }

    // return the b.msg and the error variable
    return nil
}

```

As you can see from the test output, Listing 13.45, every call to the `Build` function takes a long time to complete, and each call performs the same task.

Listing 13.45 Output Confirming the `Build` Function Runs Every Time It Is Called

```

func Test_Once(t *testing.T) {
    t.Parallel()

    b := &Builder{}

    for i := 0; i < 5; i++ {

        err := b.Build()
        if err != nil {
            t.Fatal(err)
        }

        fmt.Println("builder built")

        if !b.Built {
            t.Fatal("expected builder to be built")
        }
    }
}

```

```

$ go test -v

=== RUN   Test_Once
=== PAUSE Test_Once
=== CONT Test_Once
building..built
builder built

```

(continued)


```

building...built
builder built

building...built
builder built
building...built
builder built
building...built
builder built
--- PASS: Test_Once (0.05s)
PASS
ok      demo    0.265s

```

```

Go Version: go1.19

```

Implementing Once

As shown in Listing 13.46, you can use the `sync.Once` type inside the `Build` function to ensure that the expensive task is only performed once.

Listing 13.46 Using `sync.Once` to Run a Function Once

```

type Builder struct {
    Built bool
    once  sync.Once
}

func (b *Builder) Build() error {

    var err error

    b.once.Do(func() {

        fmt.Print("building...")

        time.Sleep(10 * time.Millisecond)

        fmt.Println("built")

        b.Built = true

        // validate the message
        if !b.Built {

```

```
        err = fmt.Errorf("expected builder to be built")
    }
})

// return the b.msg and the error variable
return err
}
```

As you can see from the test output, Listing 13.47, the `Build` function now performs the expensive task only once, and subsequent calls to the function are very fast.

Listing 13.47 Output Confirming the `Build` Function Runs Only Once

```
$ go test -v

=== RUN   Test_Once
=== PAUSE Test_Once
=== CONT Test_Once
building...built
builder built
builder built
builder built
builder built
builder built
builder built
--- PASS: Test_Once (0.01s)
PASS
ok      demo    0.248s
```

```
Go Version: go1.19
```

Closing Channels with `Once`

The `sync.Once` type is useful for closing channels. When you want to close a channel, you need to ensure that the channel is closed only once. If you try to close the channel more than once, you get a panic, and the program crashes.

Consider the example in Listing 13.48. The `Quit` method on the `Manager` is in charge of closing the `quit` channel when the `Manager` is no longer needed.

Listing 13.48 If Called Repeatedly, the `Quit` Function Panics and Closes an Already-Closed Channel

```

type Manager struct {
    quit chan struct{}
}

func (m *Manager) Quit() {
    fmt.Println("closing quit channel")
    close(m.quit)
}

```

If, however, the `Quit` method is called more than once, we are trying to close the channel more than once. We get a panic, and the program crashes.

As you can see in Listing 13.49, the tests failed as a result of trying to close the channel more than once and caused a panic.

Listing 13.49 Panicking When Trying to Close a Channel Multiple Times

```

func Test_Closing_Channels(t *testing.T) {
    t.Parallel()

    func() {
        // defer a function to catch the panic
        defer func() {

            // recover the panic
            if r := recover(); r != nil {
                // mark the test as a failure
                t.Fatal(r)
            }
        }()

        m := &Manager{
            quit: make(chan struct{}),
        }

        // close the manager's quit channel
        m.Quit()

        // try to close the manager's quit channel again
        // this will panic
        m.Quit()
    }
}

```

```

    }()
}

```

```

$ go test -v

=== RUN   Test_Closing_Channels
=== PAUSE Test_Closing_Channels
=== CONT Test_Closing_Channels
closing quit channel
closing quit channel
    demo_test.go:31: close of closed channel
--- FAIL: Test_Closing_Channels (0.00s)
FAIL
exit status 1
FAIL    demo    0.667s

```

```
Go Version: go1.19
```

In Listing 13.50, we use the `sync.Once` type to ensure that the `Quit` method, regardless of how many times it is called, only closes the channel once.

Listing 13.50 Using `sync.Once` to Close a Channel Only Once

```

type Manager struct {
    quit chan struct{}
    once sync.Once
}

func (m *Manager) Quit() {

    // close the manager's quit channel
    // this will only close the channel once
    m.once.Do(func() {
        fmt.Println("closing quit channel")
        close(m.quit)
    })
}

```

As you can see from the test output, Listing 13.51, the `Quit` method now closes the channel only once, and subsequent calls to the `Quit` method have no effect.

Listing 13.51 Output Confirming the Quit Method Closes the Channel Only Once

```
func Test_Closing_Channels(t *testing.T) {
    t.Parallel()

    m := &Manager{
        quit: make(chan struct{}),
    }

    // close the manager's quit channel
    m.Quit()

    // try to close the manager's quit channel again
    // this will now have no effect
    m.Quit()
}
```

```
$ go test -v

=== RUN   Test_Closing_Channels
=== PAUSE Test_Closing_Channels
=== CONT  Test_Closing_Channels

closing quit channel
--- PASS: Test_Closing_Channels (0.00s)
PASS
ok      demo    0.523s
```

```
Go Version: go1.19
```

Summary

In this chapter, we took a look at just a few of the synchronization types and functions in Go. First, we explored how to use a `sync.WaitGroup` to wait for a number of goroutines to finish. Then we explained how to use `sync.ErrGroup` to wait for a number of goroutines to finish and return an error if any of them failed. Next, we discussed how to use `sync.Mutex` and `sync.RWMutex` to synchronize access to a shared resource. Finally, we covered how to use `sync.Once` to ensure a function is only executed one time.

Index

SYMBOLS

<- channel operator, 344–345
%d verb, 63–66
%f verb, 66–67
//go:embed directive, 518, 521–522
\n escape sequence, 61–62
[N] syntax, 71–72
%q verb, 63
%s verb, 63
\t escape sequence, 61–62
%T verb, 67
~ (tilde) constraint operator, 325–326
%v verb, 67–69
%#v verb, 69–70
%+v verb, 169, 185
%w verb, 294, 301

A

abbreviations in package names, 6
accessing struct fields, 172–173
adding goroutines to waitgroups
 negative number of, 425–426
 positive number of, 423–424
 zero number of, 424–425
aliases, resolving name conflicts, 16, 56–57
All method, 208–209
anonymous functions, 150, 159–160
any keyword, 242–243
APIs, scope, 10
append function, 85–90, 97
 appending
 to files, 500–503
 to slices, 85–90, 97
architecture-independent numeric types, 31
arguments, 141–142
 accepting from other functions, 151
 functions as, 148–149
 return arguments, 143–144
 capturing panic values, 269–273
 error handling, 262–263
 multiple, 144–145
 named returns, 145–147
 of same type, 142–143
 variadic arguments, 151–152

 passing slices with, 153–154
 position of, 152–153
 when to use, 154–156
arrays
 avoiding out-of-bounds errors, 287–289
 converting to slices, 102–104
 creating, 79
 data types in, 82–83
 indexing, 81–82
 initializing, 79–80
 iteration over, 105
 purpose of, 77
 setting values of, 85
 slices versus, 77–79
 subsets of, 103–104
 type definitions, 83–84
 zero values, 81
AsError interface, 303–304
asserting
 interface implementation, 241–242
 with type assertions, 250–251
 avoiding panics, 284–287
 concrete type verification, 252
 generics and, 329–331
 with switch statement, 252–254
 usage example, 255–257
 validating, 251–252
assigning
 functions to variables, 283–284
 multiple values to variables, 45–46
 variables, 41–42

B

binaries, embedding files in, 520–521
bitwise operators, 135
blocking channels, 343–344
boolean logic
 boolean operators, 134
 comparison operators, 135
 equality comparison, 129–130
break keyword, 107–108
buffered channels, 358–359
 message delivery and, 360–362
 reading closed, 362–363
 usage example, 359–360
byte numeric type, 31

C

-
- callback interfaces, defining, 255–256
 - calling
 - functions without interfaces, 235–237
 - generics, 316–317
 - methods, 177
 - cancellation deadlines on contexts, 374
 - cancellation events
 - on contexts, 396
 - canceling contexts, 398–400
 - creating cancelable contexts, 397–398
 - deadlines, 405–406
 - errors on, 409–410
 - listening for, 375–376
 - listening for confirmation, 400–405
 - timeouts, 407–408
 - on error groups, listening for, 439–440
 - cap function, 92, 96–97, 112–113
 - capacity
 - of maps, 112–113
 - of slices, 91–92, 96–97
 - capitalization for exporting variables, 57
 - capturing
 - panic values, 269–273
 - system signals, 363
 - implementing graceful shutdown, 365–367
 - listening for shutdown confirmation, 368–370
 - listening for system signals, 367–368
 - os/signals package, 363–365
 - timing out nonresponsive shutdown, 370–371
 - casting values, 45
 - chan keyword, 343
 - channels
 - blocking/unblocking, 343–344
 - buffered, 358–359
 - message delivery and, 360–362
 - reading closed, 362–363
 - usage example, 359–360
 - capturing system signals, 363
 - implementing graceful shutdown, 365–367
 - listening for shutdown confirmation, 368–370
 - listening for system signals, 367–368
 - os/signals package, 363–365
 - timing out nonresponsive shutdown, 370–371
 - characteristics of, 343
 - closing, 352–354
 - closing already-closed, 357
 - reading buffered, 362–363
 - reading closed, 354–357
 - with sync.Once type, 463–466
 - writing closed, 358
 - creating, 343–344
 - listening with select statements, 348–349
 - message queues versus, 349–351
 - range loops and, 348
 - sending/receiving values, 344–345
 - unidirectional, 351–352
 - usage example, 345–347
 - choosing numeric types, 32
 - cleaning up test helpers, 227–229
 - closing channels, 352–354
 - closing already-closed, 357
 - reading buffered, 362–363
 - reading closed, 354–357
 - with sync.Once type, 463–466
 - writing closed, 358
 - closures, 149–150
 - code coverage
 - coverage profiles, 203–204
 - editor support for, 206
 - go test command, 202–203
 - go tool cover command, 204–205
 - HTML coverage reports, 205–206
 - comparison operators, 135
 - compiled languages, 28–30
 - complex values in maps, 123–124
 - updating, 125–126
 - complex64 numeric type, 31
 - complex128 numeric type, 31
 - concrete types
 - compilation errors for type assertions, 330–331
 - interfaces versus, 231–233
 - verifying via type assertions, 252
 - concurrency. *See also* synchronization
 - channels
 - blocking/unblocking, 343–344
 - buffered, 358–359
 - buffered message delivery, 360–362
 - buffered usage example, 359–360
 - capturing system signals, 363
 - characteristics of, 343
 - closing, 352–354
 - closing already-closed, 357
 - creating, 343–344
 - implementing graceful shutdown, 365–367
 - listening for shutdown confirmation, 368–370
 - listening for system signals, 367–368
 - listening with select statements, 348–349
 - message queues versus, 349–351
 - os/signals package, 363–365
 - range loops and, 348
 - reading closed, 354–357
 - reading closed buffered, 362–363
 - sending/receiving values, 344–345
 - timing out nonresponsive shutdown, 370–371
 - unidirectional, 351–352
 - usage example, 345–347
 - writing closed, 358
 - contexts
 - cancellation events, 375–376, 396–410
 - deadlines on, 374
 - default implementation, 379–380
 - error messages, 376, 408–411
 - exporting keys, 393–394
 - key security, 394–396
 - listening for system signals, 411–413
 - nodal hierarchy, 381–387

- retrieving values, 376–378
- rules for, 380–381
- string keys, 388–393
- testing system signals, 413–416
- wrapping with values, 382, 384–386
- dog feeding metaphor, 336–338
- goroutines, 338
 - memory management, 339
 - runtime.GOMAXPROCS function, 341–342
 - scheduling, 339–342
 - usage example, 342
 - work sharing, 339–340
 - work stealing, 341
- parallelism versus, 335–336
- concurrent monitors, starting, 401–402
- const keyword, 48
- constants, 48–51
 - declaring, 48
 - modifying, 48
 - typed, 49
 - untyped, 50–51
- constraints. *See* type constraints
- constraints package, 326–328
- constraints.Integer constraint, 327
- constraints.Ordered constraint, 327–328
- constraints.Signed constraint, 327
- context package, 373, 378. *See also* contexts
- context.Background function, 378–380
- context.Canceled error, 409–410
- context.CancelFunc function, 397–398, 402–403
- context.Context interface, 374, 379–380
- context.Context.Deadline method, 374
- context.Context.Done method, 375–376, 398–399
- context.Context.Err method, 376, 408–409
- context.Context.Value method, 376–378, 386–387
- context.DeadlineExceeded error, 410–411
- contexts
 - cancellation events, 396
 - canceling contexts, 398–400
 - creating cancelable contexts, 397–398
 - deadlines, 405–406
 - errors on, 409–410
 - listening for, 375–376
 - listening for confirmation, 400–405
 - timeouts, 407–408
 - deadlines on, 374
 - default implementation, 379–380
 - error messages
 - context.Canceled error, 409–410
 - context.Context.Err method, 376, 408–409
 - context.DeadlineExceeded error, 410–411
 - exporting keys, 393–394
 - key security, 394–396
 - listening for system signals, 411–413
 - nodal hierarchy, 381
 - context.WithValue function, 382, 384–386
 - key resolution, 386–387
 - usage example, 382–384
 - retrieving values, 376–378
 - rules for, 380–381
 - string keys, 388–393
 - custom types, 390–393
 - key collisions, 388–390
 - testing system signals, 413–416
 - wrapping with values, 382, 384–386
- context.WithCancel function, 397
- context.WithDeadline function, 405–406
- context.WithTimeout function, 407–408
- context.WithValue function, 382, 384–386
- continue keyword, 107–108
- control structures
 - if statements
 - else if statements, 131–132
 - else statements, 130–131
 - explained, 129–130
 - variable scope, 132–134
 - logic and math operators, 134–135
 - switch statements, 135–138
 - default blocks, 138–139
 - fallthrough keyword, 139–140
- converting
 - arrays to slices, 102–104
 - strings to/from numbers, 72–75
- copy function, 101–102
- copying
 - map values on insert, 124–125
 - slices, 101–102
- coverage profiles, 203–204
- custom data types, 167–168
- custom errors
 - defining, 291–293
 - error interface, 289
 - standard errors versus, 289–291
 - unwrapping, 298–301
- custom string key types, 390–393
- cyclical imports, 24–25

D

- data races
 - detecting in tests, 447–449
 - race flag, 445–447
 - shared access example, 443–445
- data types
 - in arrays, 82–83
 - concrete types versus interfaces, 231–233
 - constants, 49–51
 - declaring new, 167–168
 - in function arguments, 141–143
 - functions as, 147, 180–181
 - generics, 332–334
 - inheritance, 181–182
 - numeric, 31–35
 - architecture-independent, 31
 - choosing, 32
 - implementation-specific, 31

- overflow versus wraparound, 32–34
 - saturation, 34–35
- in slices, 82–83
- static versus dynamic typing, 27–28
- string literals, 35–37
 - interpreted, 35
 - quotation marks for, 35
 - raw, 36–37
- third-party, methods on, 178–179
- UTF-8 characters, 37–40
 - iterating over, 38–40
 - runes, 38
- deadlines on contexts, 374, 405–406, 410–411
- debug.PrintStack function, 307
- debug.Stack function, 307
- declaring. *See also* defining
 - constants, 48
 - data types, 167–168
 - methods, 176
 - slices, 95–96
 - variables, 40, 44–45
- decrementing integers, 122
- default blocks, 138–139
- default implementations of contexts, 379–380
- defer keyword, 156, 229, 266
- deferring function calls, 156
 - anonymous functions, 159–160
 - exit and fatal messages, 158–159
 - with multiple returns, 156–157
 - order of execution, 157
 - panics and, 158
 - scope, 160–162
 - testing.TB.Cleanup method versus, 229
- defining. *See also* declaring
 - callback interfaces, 255–256
 - custom errors, 291–293
 - interfaces, 243–247
 - structs, 168–169
 - test helpers, 222–226
 - type constraints, 321–322
- delete function, 118–119
- deleting map keys, 118–119
- dependencies, 16–22
 - go.sum file, 20–21
 - requiring, 18–20
 - updating, 21–22
 - usage, 17–18
- directories. *See also* files
 - creating, 481–483, 499–500
 - creating multiple, 489–492
 - creating structure with file path helpers, 486–487
 - errors with file path helpers, 488–489
 - fixing walk tests, 497–500
 - names, retrieving, 486
 - reading, 468–470
 - skipping, 477–481
 - special, 467–468
 - for specific files, retrieving, 485

- walking, 473–477, 510–513
- Windows filepaths, 505–506
- disabling test caching, 221
- discarding variables, 47
- do while loops, 108–109
- dog feeding metaphor (concurrency), 336–338
- dynamic typing, 27–28

E

- else if statements, 131–132
- else statements, 130–131
- embed package, 517
- embedding
 - files, 517
 - in binaries, 520–521
 - modifying, 521–522
 - as strings/byte slices, 522
 - usage example, 518–519
 - interfaces, 249–250
- embed.FS type, 518–519
- empty interfaces, 242–243
- encoding with struct tags, 173–176
- equality comparison, 129–130
- errgroup.Group type, 436, 440–442
- errgroup.Group.Go method, 437
- errgroup.Group.Wait method, 437–439
- errgroup.WithContext function, 439–440
- error groups
 - errgroup.Group type, 436, 440–442
 - errgroup.Group.Go method, 437
 - errgroup.Group.Wait method, 437–439
 - errgroup.WithContext function, 439–440
 - golang.org/x/sync/errgroup package, 434
 - problem solved by, 434–436
- error interface, 261, 289
- errors. *See also* testing
 - arrays versus slices, 103
 - in compilation, 18, 29–30, 34
 - context error messages
 - context.Canceled error, 409–410
 - context.Context.Err method, 376, 408–409
 - context.DeadlineExceeded error, 410–411
 - custom
 - defining, 291–293
 - error interface, 289
 - standard errors versus, 289–291
 - unwrapping, 298–301
 - errors.As function, 302–304
 - errors.Is function, 304–306
 - with file path helpers, 488–489
 - with formatting verbs, 70–71
 - nil receivers, 192–193
 - out of bounds, 81–82
 - panics
 - assigning functions to variables, 283–284
 - avoiding, 273
 - capturing and returning values, 269–273

- checking for nil receivers, 274–275
- checking type assertions, 284–287
- initializing interfaces, 280–283
- initializing maps, 275–278
- initializing pointers, 278–280
- log.Fatal function, 273
- from out-of-bounds errors, 264–265, 287–289
- raising, 265
- recovering from, 265–269
- returning, 144
- stack traces, 307–308
- types of, 296
- unwrapping, 297–301
- as values, 259–261
 - error interface, 261
 - errors.New function, 261
 - fmt.Errorf method, 262
 - handling, 262–263
 - usage example, 263–264
- in walk tests, fixing, 497–500
- wrapping, 294–297
- errors.As function, 302–304
- errors.Is function, 304–306, 489
- errors.New function, 261
- errors.Unwrap function, 297–301
- escape characters
 - escaping, 62–63
 - for formatting strings, 61–62
 - in raw string literals, 36
- exceptions, 259–260
- exits, deferred function calls and, 158–159
- expanding slices with variadic operator, 153–154
- explicit argument indexes, 71–72
- explicit implementation of interfaces, 233–234
- exporting
 - context keys, 393–394
 - variables, 57
- extensions, retrieving, 485

F

- failing fast when testing, 220–221
- fallthrough keyword, 139–140
- fatal message, deferred function calls and, 158–159
- file systems
 - fs package, 506–507
 - fs.File interface, 509
 - fs.FS interface, 508, 510–513
 - mocking, 513–517
- file tree, example of, 467–468
- filepath package, 484
- filepath.Base function, 486
- filepath.Dir function, 485, 486–487
- filepath.Ext function, 485, 486–487
- filepath.Join function, 505–506
- filepath.WalkDir function, 473–474, 475–477
- files. *See also* directories
 - appending to, 500–503

- creating, 492–495, 499–500
- directory of, retrieving, 485
- embedding, 517
 - in binaries, 520–521
 - modifying, 521–522
- as strings/byte slices, 522
- usage example, 518–519
- errors with file path helpers, 488–489
- extensions, retrieving, 485
- fixing walk tests, 497–500
- metadata on, 471–473
- names, retrieving, 486
- naming, 9
- reading, 494–495, 503–505
- skipping, 477
- truncating, 495–497
 - Windows filepaths, 505–506
- first-class functions, 147
- float32 numeric type, 31–32
- float64 numeric type, 31–32
- floats, formatting with formatting verbs, 66–67
- fmt package, functions in, 57–58
 - errors with formatting verbs, 70–71
 - escape characters, 61–63
 - explicit argument indexes, 71–72
 - float formatting verbs, 66–67
 - formatting verbs, 61
 - Fprint functions, 59
 - integer formatting verbs, 63–66
 - multiple Printf arguments, 61
 - new lines, 60
 - Print functions, 59
 - Sprint functions, 58–59
 - string formatting verbs, 63
 - value Go-syntax printing, 69–70
 - value printing, 67–69
 - value type printing, 67
- fmt.Errorf method, 262, 294, 301
- fmt.Stringer interface, 241
- folders, package folder structure, 4, 6–13. *See also* directories
 - multiple packages, 7–9
 - naming files, 9
 - organization, 9–13
- for loops
 - appending to slices, 88–89
 - break keyword, 107–108
 - continue keyword, 107–108
 - explained, 104–105
 - iteration over arrays and slices, 105
 - range keyword, 105–106
- formatting
 - floats with formatting verbs, 66–67
 - integers with formatting verbs, 63–66
 - strings
 - errors with formatting verbs, 70–71
 - escape characters, 61–63
 - explicit argument indexes, 71–72
 - with formatting verbs, 61, 63
 - list of functions, 57–58
 - Sprint functions, 58–59

- verbs for, 61
 - errors with, 70–71
 - escape characters, 61–63
 - explicit argument indexes, 71–72
 - floats, 66–67
 - integers, 63–66
 - strings, 63
 - value type printing, 67
- Fprint functions, 59
- fs package, 506–507
- fs.DirEntry type, 469–470
- fs.ErrExit error, 489
- fs.File interface, 509
- fs.FileInfo interface, 471–472
- fs.FS interface, 508, 510–513
- fs.SkipDir error, 478–479
- fstest.MapFile type, 509, 515
- fstest.MapFS type, 514, 515–517
- fs.WalkDir function, 510–512
- fs.WalkDirFunc function, 474–475
- functions. *See also* interfaces
 - anonymous, 150
 - arguments, 141–142
 - accepting from other functions, 151
 - capturing panic values, 269–273
 - error handling, 262–263
 - functions as, 148–149
 - multiple return arguments, 144–145
 - named returns, 145–147
 - passing slices with, 153–154
 - position of variadic arguments, 152–153
 - return arguments, 143–144
 - of same type, 142–143
 - variadic arguments, 151–152
 - when to use variadic arguments, 154–156
- assigning to variables, 283–284
- calling without interfaces, 235–237
- closures, 149–150
- as data types, 180–181
- deferring calls, 156
 - anonymous functions, 159–160
 - exit and fatal messages, 158–159
 - with multiple returns, 156–157
 - order of execution, 157
 - panics and, 158
 - scope, 160–162
 - testing.TB.Cleanup method versus, 229
- definitions, 141
- as first class, 147
- in fmt package, 57–58
 - errors with formatting verbs, 70–71
 - escape characters, 61–63
 - explicit argument indexes, 71–72
 - float formatting verbs, 66–67
 - formatting verbs, 61
 - Fprint functions, 59
 - integer formatting verbs, 63–66
 - multiple Println arguments, 61
 - new lines, 60

- Print functions, 59
- Sprint functions, 58–59
- string formatting verbs, 63
 - value Go-syntax printing, 69–70
 - value printing, 67–69
 - value type printing, 67
- generics, 311
 - calling, 316–317
 - constraints package, 326–328
 - defining constraints, 321–322
 - instantiating, 320–321
 - mixing type and method constraints, 331
 - multiple type constraints, 322–323
 - multiple types, 317–320
 - type assertions, 329–331
 - type constraints, 315–317
 - underlying type constraints, 324–326
- init, 162–163
 - multiple, 163–164
 - order of execution, 164
 - for side effects, 164–166
- methods on, 181
- methods versus, 177
- new, 189–190
- in strconv package, 72–75
- test function signatures, 196
- test helpers, 222
 - cleaning up, 227–229
 - defining, 222–226
 - marking functions as, 226–227
- as variables, 147–148

G

- garbage collection, 28
- generic functions, 311
 - calling, 316–317
 - constraints package, 326–328
 - defining constraints, 321–322
 - instantiating, 320–321
 - mixing type and method constraints, 331
 - multiple type constraints, 322–323
 - multiple types, 317–320
 - type assertions, 329–331
 - type constraints, 315–317
 - underlying type constraints, 324–326
- generic package names, avoiding, 6
- generic types, 332–334
- getting pointers, 185–186
- go build command, 29
- go get command, 4, 18–20
- go get -u command, 21–22
- go help mod command, 2
- go help test command, 477
- go mod command, 2
- go mod init command, 3
- go test command, 202–203

go tool cover command, 204–205
 go vet command, 71
 golang.org/x/exp/constraints package, 326–328
 golang.org/x/sync/errgroup package, 434
 go.mod file, 2
 goroutines, 338. *See also* channels; contexts; synchronization
 memory management, 339
 messages and, 350–351
 runtime.GOMAXPROCS function, 341–342
 scheduling, 339–342
 usage example, 342
 work sharing, 339–340
 work stealing, 341
 go.sum file, 20–21
 Go-syntax of value, printing, 69–70
 graceful shutdown
 implementing, 365–367
 listening for shutdown confirmation, 368–370
 listening for system signals, 367–368
 rules for, 363
 timing out nonresponsive shutdown, 370–371
 growing slices, 92–94, 97–99

H

handling errors, 262–263
 heap, 183
 helper functions
 in context package, 378
 test helpers, 222
 cleaning up, 227–229
 defining, 222–226
 marking functions as, 226–227
 .hidden directory, 467–468
 HTML coverage reports, 205–206
 http.HandlerFunc type, 181

I

identifiers. *See* constants; variables
 if statements, 129–130
 else if statements, 131–132
 else statements, 130–131
 variable scope, 132–134
 _ignore directory, 467–468
 implementation-specific numeric types, 31
 implicit implementation of interfaces, 234–235
 import aliases, 16
 import keyword, 14
 importing
 modules, 13–16, 22
 packages, 13–16
 cyclical imports, 24–25
 multiple versions, 23–24
 for side effects, 164–166

incrementing integers, 122
 indexing
 arrays, 81–82
 avoiding out-of-bounds errors, 287–289
 slices, 81–82
 inferred typing, 50–51
 inheritance, 181–182
 init functions, 162–163
 multiple, 163–164
 order of execution, 164
 for side effects, 164–166
 initializing
 arrays, 79–80
 interfaces, 280–283
 maps, 113, 275–278
 modules, 3
 packages. *See* init functions
 pointers, 278–280
 slices, 79–80
 structs, 169–172
 variables, 44–45
 instantiating generics, 320–321
 int numeric type, 31–32
 int8 numeric type, 31
 int16 numeric type, 31
 int32 numeric type, 31
 int64 numeric type, 31
 integers
 formatting with formatting verbs, 63–66
 incrementing/decrementing, 122
 padding, 64–66
 interfaces
 AsError interface, 303–304
 asserting implementation, 241–242
 callback, 255–256
 calling functions without, 235–237
 concrete types versus, 231–233
 context.Context interface, 374
 defining, 243–247
 drawbacks of, 311–315
 embedding, 249–250
 empty, 242–243
 error interface, 261, 289
 explicit implementation, 233–234
 implementing, 247–249
 implicit implementation, 234–235
 initializing, 280–283
 io.Writer interface, 237–240
 IsError interface, 305–306
 multiple, 241
 type assertions, 250–251
 concrete type verification, 252
 with switch statement, 252–254
 usage example, 255–257
 validating, 251–252
 Unwrapper interface, 299
 usage example, 237–239
 validatable, 250
 interpreted languages, 28–29

interpreted string literals, 35
 io.Reader interface, 503–505
 io.ReadWriteCloser interface, 249–250
 io.Writer interface, 237–240, 503–505
 IsError interface, 305–306
 iterating over UTF-8 characters, 38–40
 iteration

- of arrays and slices, 105
- do while loops, 108–109
- for loops
 - appending to slices, 88–89
 - break keyword, 107–108
 - continue keyword, 107–108
 - explained, 104–105
 - range keyword, 105–106
- of maps, 116–118

J

JSON encoding with struct tags, 173–176

K

keys

- in contexts
 - collisions, 388–390
 - custom string key types, 390–393
 - exporting, 393–394
 - resolving, 386–387
 - security, 394–396
 - string keys, 388
- in maps, 114–116
 - deleting, 118–119
 - interface drawbacks, 311–315
 - listing, 126–127
 - nonexistent, 120–121
 - sorting, 127–129
 - testing presence of, 122–123

keywords, list of, 30

L

len function, 91, 96–97, 112

length

- of maps, 112–113
- of slices, 91–92, 96–97

listener function, 353

listening

- for cancellation events on contexts, 375–376
- to channels
 - with select statements, 348–349
 - for shutdown confirmation, 368–370
 - for system signals, 367–368

- for context cancellation confirmation, 400–405

- for error group cancellation, 439–440

- for system signals with contexts, 411–413

listing map keys, 126–127

listings

- accessing struct fields, 172–173

- adding negative number of goroutines to waitgroups, 425–426

- adding positive number of goroutines to waitgroups, 423–424

- adding zero number of goroutines to waitgroups, 424–425

a.go file, 8

All method, 208–209

anonymous functions, 150

any keyword, 242–243

appending to files, 501–503

appending to slices, 86–90

array of four strings, 78

AsError interface, 303–304

asserting interface implementation, 242

assigning functions to variables, 283–284

assigning multiple values, 45–46

assigning variables, 41–42

bad folder files listing, 8

bad identifier names, 52

bad package build, 9

bad.go file, 10–11

bar/bar.go file, 15

bar.go file, 24–25

b.go file, 8

blocking/unblocking channels, 345

boolean equality comparison, 129–130

break keyword, 107–108

buffered channel message delivery, 361–362

buffered channel usage example, 359–360

calling functions without interfaces, 235–237

calling generics, 316–317

calling methods, 177

cap function, 92, 112–113

capturing and returning panic values, 269–273

casting values, 45

channel usage example, 345–347

checking for nil receivers, 193, 274–275

checking type assertions, 285–287

closing already-closed channels, 357

closing channels, 353–354, 463–466

closures, 149–150

code coverage, 203

compilation errors, 18, 29–30, 34

complex if/else if statements, 135–136

complex values in maps, 123–124

concrete types versus interfaces, 231–233

constraints.Integer constraint, 327

constraints.Ordered constraint, 328

constraints.Signed constraint, 327

context key collisions, 389–390

context key security, 394–396

context package, 373, 378

- context.Background function, 378–380
- context.Canceled error, 409–410
- context.CancelFunc function, 397–398, 402–403
- context.Context interface, 374
- context.Context.Deadline method, 374
- context.Context.Done method, 375–376, 398–399
- context.Context.Err method, 376, 408–409
- context.Context.Value method, 376–378
- context.DeadlineExceeded error, 410–411
- context.WithCancel function, 397
- context.WithDeadline function, 405–406
- context.WithTimeout function, 407–408
- context.WithValue function, 385–386
- continue keyword, 107–108
- converting arrays to slices, 102–104
- converting strings to/from numbers, 74–75
- copying map values on insert, 124–125
- copying slices, 101–102
- coverage profiles, 203–204
- coverage profiles on per-function basis, 207
- crafting failure messages, 201–202
- creating arrays versus slices, 79
- creating directory structure, 486–487
- creating files/directories, 499–500
- custom data type usage, 168
- custom string key types, 391–393
- cyclical imports, 24–25
- data types in arrays and slices, 83
- debug.PrintStack function, 307
- debug.Stack function, 307
- declaring constants, 48
- declaring methods, 176
- declaring new data types, 167
- declaring slices, 95–96
- declaring variables, 40
- default blocks in switch statements, 138–139
- deferred function calls and exit/fatal message, 159
- deferred function calls and panics, 158
- deferring anonymous functions, 159–160
- deferring function calls, 156
- deferring function calls with multiple returns, 156–157
- defining custom errors, 292–293
- defining interfaces, 243–247
- defining structs, 169
- defining test helpers, 223–226
- defining type constraints, 321–322
- delete function, 118
- deleting map keys, 118–119
- demo module files listing, 15
- detecting race conditions in tests, 447–449
- disabling test caching, 221
- discarding variables, 47
- do while loops, 109
- else if statements, 132
- else statements, 130–131
- embed package, 517
- embedding files as strings/byte slices, 522
- embedding files in binaries, 520–521
- embed.FS type, 518–519
- empty interface, 242, 243
- encoding with struct tags, 173–175
- errgroup.Group type, 436, 441–442
- errgroup.Group.Go method, 437
- errgroup.Group.Wait method, 437–439
- errgroup.WithContext function, 439–440
- error interface, 261, 289
- error types, 296
- error usage example, 263–264
- errors as values, 260–261
- errors with file path helpers, 488–489
- errors with formatting verbs, 70
- errors.As function, 302–303
- errors.Is function, 305, 489
- errors.New function, 261
- errors.Unwrap function, 297–301
- escape characters \n and \t, 62
- escaping escape characters, 62–63
- exceptions, 259–260
- expanding slices with variadic operator, 154
- explicit argument indexes, 71–72
- explicit implementation of interfaces, 233–234
- exporting context keys, 393–394
- exporting variables, 57
- external dependencies file structure, 20
- failing fast when testing, 221
- fallthrough keyword, 139–140
- file tree example, 467–468
- filepath package, 484
- filepath.Base function, 486
- filepath.Dir function, 485
- filepath.Ext function, 485
- filepath.Join function, 505–506
- filepath.WalkDir function, 473–474, 475–477
- fixing walk tests, 497–500
- flect.Dasherize function, 17
- float formatting verbs, 66
- fint package, 58
- fint.Errorf method, 262
- fint.Fprint function, 59
- fint.Print function, 59
- fint.Println function, 60
- fint.Sprintf function, 58–59
- fint.Stringer interface, 241
- foo type, 25
- foo/bar/bar.go file, 15
- foo.go file, 13
- for loops, 105
- fs package, 507
- fs.DirEntry type, 469–470
- fs.File interface, 509
- fs.FileInfo interface, 471–472
- fs.FS interface, 508
- fs.SkipDir error, 478–479
- fstest.MapFile type, 509, 515
- fstest.MapFS type, 514, 515–517
- fs.WalkDir function, 510–512
- fs.WalkDirFunc function, 474–475
- function arguments, 142–143

- function definitions, 141
- function types, 180–181
- functions accepting other function arguments, 151
- functions as arguments, 148–149
- functions as variables, 147–148
- generic types, 332–334
- generics, 316
- getting pointers, 185–186
- //go:embed directive, 518, 521–522
- go get command, 4, 18–19
- go help test command, 477
- go mod command, 2
- go mod init command, 3
- Go module files listing, 1
- go tool cover command, 204–205
- go vet command, 71
- golang.org/x/exp/constraints package, 326
- golang.org/x/sync/errgroup package, 434
- go.mod file, 2, 18
- go.mod file updates, 20
- good folder files listing, 7–8, 11–12
- goroutine usage example, 342
- goroutines, 338
- go.sum file, 21
- growing slices, 92–94
- handling errors, 263
- HTML coverage reports, 205–206
- http.HandlerFunc type, 181
- if statements, 130
- implementing graceful shutdown, 365–367
- implementing interfaces, 247–249
- implicit implementation of interfaces, 234–235
- importing modules with semantic versioning, 22
- importing multiple packages, 14, 23–24
- improper mutex usage, 456–459
- indexing arrays and slices, 81–82
- inferred typing, 50–51
- inheritance, 182
- init function, 162–163
- init functions for side effects, 165–166
- initializing arrays and slices, 80
- initializing interfaces, 280–283
- initializing maps, 113, 276–278
- initializing pointers, 278–280
- initializing structs, 169–172
- initializing variables, 44
- instantiating generics, 320–321
- integer formatting verbs, 64
- interface usage example, 238–239
- interpreted string literals, 35
- io.ReadWriteCloser interface, 249–250
- io.Writer interface, 237, 239–240
- IsError interface, 305–306
- iterating over maps, 116–118
- iterating over UTF-8 characters, 39–40
- key resolution in contexts, 386–387
- len function, 91, 112
- listener function, 353
- listening for shutdown confirmation, 368–370
- listening for system signals, 367–368, 412–413
- listing map keys, 126–127
- log.Fatal function, 273
- logging in tests, 215
- main.go file, 14, 17
- make function, 95, 113, 344
- make function with append function, 97
- make function with length and capacity, 96–97
- managing goroutines with sync.WaitGroup, 434–436
- map keys, 114–115, 311–315
- messages and goroutines, 350–351
- method receivers, 188–189
- methods on third-party types, 178–179
- mixing type and method constraints, 331
- modifying constants, 48
- modifying variables, 48–49
- multiple generic types, 317–319
- multiple goroutines for one task, 419–420
- multiple init functions, 163–164
- multiple interfaces, 241
- multiple Println arguments, 61
- multiple return arguments, 144–145
- multiple type constraints, 322–323
- mutating subsets of slices, 100
- name conflicts between variables and packages, 54–55
- name conflicts when importing packages, 15–16
- named returns, 146–147
- names folder files listing, 13
- naming files, 9
- naming packages, 5–6
- naming variables, 53
- new function, 189–190
- new lines, 60
- nil receivers, 191–192
- nodal hierarchy for contexts, 381–384
- nonexistent map keys, 120–121
- order of execution for deferred function calls, 157
- order of execution for init functions, 164
- os.Create function, 492–494, 495–496
- os.DirEntry type, 468
- os.DirFS function, 512
- os.FileMode type, 482
- os.Interrupt type, 365
- os.Mkdir function, 481–483
- os.MkdirAll function, 490–492
- os.OpenFile function, 501
- os.ReadDir function, 468, 470
- os.ReadFile function, 494–495
- os.Signal type, 364
- os.Stat function, 472–473
- overflow example, 33–34
- package folder structure, 4, 6–7
- padding integers, 64–66
- panic from out-of-bounds error, 264–265, 287–289
- panic function, 265
- passing by value, 184
- path package, 54
- pointer usage, 187
- printing stack traces, 308

- printing value Go-syntax, 69–70
 - printing value type, 67
 - printing values, 68–69
 - problem solved by single task execution, 460–462
 - race flag, 445–446
 - range keyword, 106
 - range loops, 348
 - raw string literals, 36–37
 - reading closed buffered channels, 362–363
 - reading closed channels, 354–355
 - reading files, 504–505
 - receiving pointers, 184
 - recover function, 266
 - recovering from panics, 266–269
 - resolving name conflicts, 55–57
 - return arguments, 143
 - returning errors, 144
 - runes, 38
 - running package tests, 213
 - running specific tests, 218–219
 - running tests with subpackages, 213–214
 - runtime.GOMAXPROCS function, 341–342
 - saturation example, 34–35
 - scope of deferred functions, 160–162
 - select statements, 349
 - sending TEST_SIGNAL signal, 416
 - sending/receiving values via channels, 345
 - setting and retrieving map values, 111–112
 - shared access example, 443–445, 449
 - signal.Notify function, 364
 - signal.NotifyContext function, 412
 - slice of four strings, 78–79
 - smtp package, 7
 - sort package, 127–128
 - sorting map keys, 128–129
 - sound package, 7
 - standard error examples, 290–291
 - starting concurrent monitors, 401–402
 - strconv.ParseInt function, 73
 - string formatting verbs, 63
 - string keys in contexts, 388
 - struct tags, 173
 - structs as map keys, 115–116
 - structure of subtests, 211–212
 - structure of table driven tests, 208
 - subsets of slices, 99
 - subtests, 212–213
 - switch statements, 136–138
 - sync package, 447
 - sync.Locker interface, 450
 - sync.Locker.Lock method, 450
 - sync.Locker.Unlock method, 451
 - sync.Mutex type, 451–453
 - sync.Once type, 459–460, 462–463, 465
 - sync.RWMutex type, 453–455
 - sync.Waitgroup type, 421, 432–433
 - sync.WaitGroup.Add method, 422
 - sync.WaitGroup.Done method, 427–431
 - sync.WaitGroup.Wait method, 422
 - syntactic sugar for methods, 177–178
 - table driven tests, 209–210
 - test file naming conventions, 196
 - test function signatures, 196
 - testing map key presence, 122–123
 - testing system signals, 414–416
 - testing.Short function, 216–217
 - testing.T type, 197
 - testing.TB interface, 222–223
 - testing.TB.Cleanup method, 227–229
 - testing.TB.Helper method, 226
 - testing.T.Error method, 198–200
 - testing.T.Fatal method, 198, 200–201
 - testing.T.Parallel method, 217–218
 - testing.T.Run method, 211
 - theoretical representation of slices, 91
 - time.Ticker function, 351
 - timing out nonresponsive shutdown, 370–371
 - timing out tests, 219–220
 - type assertion usage example, 255–257
 - type assertions, 251, 329–330
 - type assertions with switch statements, 253–254
 - type definitions for arrays and slices, 83–84
 - typed constants, 49
 - underlying type constraints, 324–326
 - unidirectional channels, 352
 - uninitialized maps, 114
 - unused variables, 47
 - Unwrapper interface, 299
 - unwrapping errors, 297–301
 - updating complex values in maps, 125–126
 - updating dependencies, 21–22
 - UTF-8 character example, 37–38
 - validating type assertions, 252
 - variable declarations, 27–28
 - variable scope in if statements, 132–134
 - variadic arguments, 151–153
 - verbose test output, 214
 - verifying concrete types with type assertions, 252
 - when to use variadic arguments, 154–156
 - work sharing, 339–340
 - work stealing, 341
 - wraparound example, 33
 - wrapping errors, 294–297
 - writing to closed channels, 358
 - zero values for arrays and slices, 81
 - zero values for closed channels, 355–357
 - zero values for maps, 121–122
 - zero values for variables, 42–44
- lists
- arrays
 - converting to slices, 102–104
 - creating, 79
 - data types in, 82–83
 - indexing, 81–82
 - initializing, 79–80
 - iteration over, 105
 - purpose of, 77
 - setting values of, 85

- slices versus, 77–79
 - subsets of, 103–104
 - type definitions, 83–84
 - zero values, 81
 - slices
 - appending to, 85–90, 97
 - arrays versus, 77–79
 - converting arrays to, 102–104
 - copying, 101–102
 - creating, 79, 95–97
 - data types in, 82–83
 - declaring, 95–96
 - growing, 92–94, 97–99
 - indexing, 81–82
 - initializing, 79–80
 - iteration over, 105
 - length and capacity, 91–92, 96–97
 - purpose of, 77
 - setting values of, 85
 - subsets of, 99–100
 - theoretical representation of, 91
 - type definitions, 83–84
 - zero values, 81
 - locking resources
 - improper mutex usage, 455–459
 - sync.Locker interface, 449–451
 - sync.Mutex type, 451–453
 - sync.RWMutex type, 453–455
 - log.Fatal function, 273
 - logging in tests, 215
 - logic. *See* boolean logic; control structures
 - loops
 - do while loops, 108–109
 - for loops
 - appending to slices, 88–89
 - break keyword, 107–108
 - continue keyword, 107–108
 - explained, 104–105
 - iteration over arrays and slices, 105
 - range keyword, 105–106
 - over channels, 348
-
- ## M
-
- make function, 95–97, 113, 278, 344
 - maps
 - explained, 111
 - initializing, 113, 275–278
 - iteration over, 116–118
 - keys in, 114–116
 - deleting, 118–119
 - interface drawbacks, 311–315
 - listing, 126–127
 - nonexistent, 120–121
 - sorting, 127–129
 - testing presence of, 122–123
 - length and capacity, 112–113
 - uninitialized, 114
 - values
 - complex values, 123–124
 - copying on insert, 124–125
 - updating complex values, 125–126
 - zero values, 121–122
 - mathematical operators, 134
 - memory management, 28
 - array and slice memory spaces, 85
 - cleaning up test helpers, 227–229
 - goroutines, 339
 - heap, 183
 - pointers and, 190–191
 - stack, 183, 186
 - message delivery, buffered channels and, 360–362
 - message queues, channels versus, 349–351
 - metadata on files, 471–473
 - methods
 - calling, 177
 - declaring, 176
 - on functions, 181
 - functions versus, 177
 - inheritance, 181–182
 - mixing type and method constraints, 331
 - syntactic sugar, 177–178
 - for test failures, 198–202
 - on third-party types, 178–179
 - value versus pointer receivers, 188–189
 - mocking file systems, 513–517
 - modifying
 - constants, 48
 - embedded files, 521–522
 - variables, 48–49
 - modules, 1–4
 - commands for, 2
 - definition of, 1
 - dependencies, 16–22
 - go.sum file, 20–21
 - requiring, 18–20
 - updating, 21–22
 - usage, 17–18
 - importing, 13–16, 22
 - initializing, 3
 - VCS and, 3–4
 - monitors, starting concurrent, 401–402
 - multiline strings in raw string literals, 36–37
 - multiple arguments with Println function, 61
 - multiple directories, creating, 489–492
 - multiple generic types, 317–320
 - multiple init functions, 163–164
 - multiple interfaces, 241
 - multiple packages
 - in folder structure, 198
 - importing, 14, 23–24
 - multiple return arguments, 144–145
 - multiple type constraints, 322–323
 - multiple values, assigning to variables, 45–46
 - mutating subsets of slices, 100
 - mutexes, synchronization with
 - detecting race conditions in tests, 447–449
 - improper usage, 455–459

locking resources, 449–451
 sync.RWMutex type, 453–455
 usage example, 451–453

N

name conflicts
 between variables and packages, 53–57
 when importing packages, 14–16
 named returns, 145–147
 names of files/directories, retrieving, 486
 naming
 files, 9
 packages, 5–6
 test files, 195–196
 variables, 51–57
 bad examples, 52
 capitalization, 57
 package name conflicts, 53–57
 rules for, 51–52
 styles for, 53
 negative number of goroutines, adding to waitgroups, 425–426
 new function, 189–190
 new lines
 escape characters for, 61–62
 printing, 60
 nil receivers, 191–193, 274–275
 nil value for variables, 42
 nodal hierarchy for contexts, 381
 context.WithValue function, 382, 384–386
 key resolution, 386–387
 usage example, 382–384
 nonexistent map keys, 120–121
 nonresponsive shutdown, timing out, 370–371
 numbers, converting strings to/from, 72–75
 numeric types, 31–35
 architecture-independent, 31
 choosing, 32
 implementation-specific, 31
 overflow versus wraparound, 32–34
 saturation, 34–35

O

operators
 list of, 30
 logic and math, 134–135
 order of execution
 for deferred function calls, 157
 for init functions, 164
 os.Create function, 492–494, 495–496, 499–500
 os.DirEntry type, 468
 os.DirFS function, 512
 os.FileMode type, 482
 os.Interrupt type, 365

os.Mkdir function, 481–483, 488–489
 os.MkdirAll function, 489–492, 499–500
 os.OpenFile function, 501
 os.ReadDir function, 468, 470
 os.ReadFile function, 494–495
 os.Signal type, 364
 os/signal package, 363–365
 os.Stat function, 472–473
 out-of-bounds errors, 81–82, 264–265, 287–289
 overflow in numeric types, 32–34

P

package aliases, resolving name conflicts, 16, 56–57
 package keyword, 5
 packages, 4–6
 definition of, 4
 dependencies, 16–22
 go.sum file, 20–21
 requiring, 18–20
 updating, 21–22
 usage, 17–18
 folder structure, 4, 6–13
 multiple packages, 7–9
 naming files, 9
 organization, 9–13
 importing, 13–16
 cyclical imports, 24–25
 multiple versions, 23–24
 for side effects, 164–166
 initializing. *See* init functions
 name conflicts
 with variables, 53–57
 when importing, 14–16
 naming, 5–6
 running tests, 213, 217
 padding integers, 64–66
 panic function, 265
 panics
 avoiding, 273
 assigning functions to variables, 283–284
 checking for nil receivers, 274–275
 checking type assertions, 284–287
 initializing interfaces, 280–283
 initializing maps, 275–278
 initializing pointers, 278–280
 from out-of-bounds errors, 287–289
 capturing and returning values, 269–273
 closing channels, 464–465
 decrementing sync.WaitGroup, 431
 deferred function calls and, 158
 log.Fatal function, 273
 from out-of-bounds errors, 264–265, 287–289
 raising, 265
 recovering from, 265–269
 parallelism
 concurrency versus, 335–336
 running tests in parallel, 217–218

ParseInt function, 73

passing

- by reference, 186
- by value, 183–184

path for importing packages, 13–14

performance, pointers and, 190–191

pointer receivers, value receivers versus, 188–189

pointers, 182–183

- getting, 185–186
- initializing, 278–280
- new function, 189–190
- nil receivers, 191–193
- passing by reference, 186
- passing by value, 183–184
- performance, 190–191
- printing, 185
- receiving, 184
- usage example, 186–187

positive number of goroutines, adding to waitgroups, 423–424

Print functions, 59

printing

- pointers, 185
- runes, 38
- stack traces, 308
- strings
 - Fprint functions, 59
 - multiple Println arguments, 61
 - new lines, 60
 - Print functions, 59
- structs, 169
- values
 - with formatting verbs, 67–69
 - Go-syntax of value, 69–70
 - type of value, 67

Printf function

- adding new lines, 60
- multiple arguments, 61

Q

quotation marks for string literals, 35

R

race conditions

- detecting in tests, 447–449
- race flag, 445–447
- shared access example, 443–445
- race flag, 445–447

raising panics, 265

range keyword, 39–40, 105–106, 116–118, 348

raw string literals, 36–37

reading

- closed buffered channels, 362–363
- closed channels, 354–357

directories, 468–470

files, 494–495, 503–505

receivers

- nil, 191–193
- value versus pointer, 188–189

receiving

- pointers, 184
- values via channels, 344–345

recover function, 265–269

recovering from panics, 265–269

requiring dependencies, 18–20

return arguments, 143–144

- accepting in functions, 151
- capturing panic values, 269–273
- error handling, 262–263
- multiple, 144–145
- named returns, 145–147

rune numeric type, 31, 38

running tests, 213

- disabling test caching, 221
- failing fast, 220–221
- logging in tests, 215
- package tests, 213, 217
- in parallel, 217–218
- short tests, 216–217
- specific tests, 218–219
- with subpackages, 213–214
- timing out, 219–220
- verbose output, 214

runtime.GOMAXPROCS function, 341–342

S

saturation in numeric types, 34–35

scheduling goroutines, 339–342

scope

- of APIs, 10
- of deferred functions, 160–162
- of variables in if statements, 132–134

security of context keys, 394–396

select statements, 348–349

selecting numeric types, 32

semantic versioning, 22

sending

- TEST_SIGNAL signal, 416
- values via channels, 344–345

shared access

- improper mutex usage, 455–459
- locking resources, 449–451
- mutex usage example, 451–453
- race conditions, 443–445, 447–449
- sync.RWMutex type, 453–455

short tests, 216–217

shutdown, graceful

- implementing, 365–367
- listening for shutdown confirmation, 368–370
- listening for system signals, 367–368
- rules for, 363
- timing out nonresponsive shutdown, 370–371

- shutdown confirmation, listening for, 368–370
- side effects, init functions for, 164–166
- signal.Notify function, 364
- signal.NotifyContext function, 412
- single execution of task
 - closing channels, 463–466
 - problem solved by, 460–462
 - sync.Once type, 459–460, 462–463
- skipping
 - directories, 477–481
 - files, 477
- slices
 - appending to, 85–90, 97
 - arrays versus, 77–79
 - avoiding out-of-bounds errors, 287–289
 - converting arrays to, 102–104
 - copying, 101–102
 - creating, 79, 95–97
 - data types in, 82–83
 - declaring, 95–96
 - embedding files as, 522
 - expanding with variadic operator, 153–154
 - growing, 92–94, 97–99
 - indexing, 81–82
 - initializing, 79–80
 - iteration over, 105
 - length and capacity, 91–92, 96–97
 - purpose of, 77
 - setting values of, 85
 - subsets of, 99–100
 - theoretical representation of, 91
 - type definitions, 83–84
 - zero values, 81
- sort package, 127–129
- sorting map keys, 127–129
- specific tests, running, 218–219
- Sprint functions, 58–59
- stack, 183, 186
- stack traces, 307–308
- standard errors, custom errors versus, 289–291
- starting concurrent monitors, 401–402
- static typing, 27–28
- strconv package, 72–75
- string keys in contexts, 388
 - custom types, 390–393
 - key collisions, 388–390
- string literals, 35–37
 - converting to/from numbers, 72–75
 - embedding files as, 522
 - formatting
 - errors with formatting verbs, 70–71
 - escape characters, 61–63
 - explicit argument indexes, 71–72
 - with formatting verbs, 61, 63
 - list of functions, 57–58
 - Sprint functions, 58–59
 - interpreted, 35
 - printing
 - Fprint functions, 59
 - multiple Println arguments, 61
 - new lines, 60
 - Print functions, 59
 - quotation marks for, 35
 - raw, 36–37
- struct tags, encoding with, 173–176
- structs, 167
 - accessing fields, 172–173
 - defining, 168–169
 - initializing, 169–172
 - as map keys, 115–116
 - printing, 169
- structural typing, 235
- subdirectories, creating, 481–483
- subpackages, running tests with, 213–214
- subsets
 - of arrays, 103–104
 - of slices, 99–100
- subtests, 211–213
- switch statements, 135–138
 - default blocks, 138–139
 - fallthrough keyword, 139–140
 - type assertions with, 252–254
- sync package, 447
- synchronization
 - error groups
 - errgroup.Group type, 436, 440–442
 - errgroup.Group.Go method, 437
 - errgroup.Group.Wait method, 437–439
 - errgroup.WithContext function, 439–440
 - golang.org/x/sync/errgroup package, 434
 - problem solved by, 434–436
 - with mutexes
 - detecting race conditions in tests, 447–449
 - improper usage, 455–459
 - locking resources, 449–451
 - sync.RWMutex type, 453–455
 - usage example, 451–453
 - race conditions
 - race flag, 445–447
 - shared access example, 443–445
 - single execution of task closing channels, 463–466
 - problem solved by, 460–462
 - sync.Once type, 459–460, 462–463
- sync package, 447
- waitgroups
 - problem solved by, 419–421
 - sync.Waitgroup type, 421, 432–433
 - sync.WaitGroup.Add method, 422–426
 - sync.WaitGroup.Done method, 426–431
 - sync.WaitGroup.Wait method, 421–422
- sync.Locker interface, 449–451
- sync.Locker.Lock method, 450
- sync.Locker.Unlock method, 451
- sync.Mutex type, 451–453
- sync.Once type, 459–460, 462–463, 465

- sync.RWMutex type, 453–455
- sync.Waitgroup type, 421, 432–433
- sync.WaitGroup.Add method, 422–426
- sync.WaitGroup.Done method, 426–431
- sync.WaitGroup.Wait method, 421–422
- syntactic sugar for methods, 177–178
- system signals
 - capturing, 363
 - implementing graceful shutdown, 365–367
 - listening for shutdown confirmation, 368–370
 - listening for system signals, 367–368
 - os/signals package, 363–365
 - timing out nonresponsive shutdown, 370–371
 - listening with contexts, 411–413
- testing, 413–416

T

- tab character, escape characters for, 61–62
- table driven testing, 206–207
 - structure of, 207–208
 - subtests, 211–213
 - writing, 208–210
- test caching, disabling, 221
- test failures
 - crafting failure messages, 201–202
 - marking, 198
 - testing.T.Error method, 198–200
 - testing.T.Fatal method, 200–201
- test files, naming, 195–196
- test function signatures, 196
- test helpers, 222
 - cleaning up, 227–229
 - defining, 222–226
 - marking functions as, 226–227
- TEST_SIGNAL signal, sending, 416
- testdata directory, 467–468, 480–481
- testing
 - map key presence, 122–123
 - system signals, 413–416
- testing framework, 195
 - code coverage
 - coverage profiles, 203–204
 - editor support for, 206
 - go test command, 202–203
 - go tool cover command, 204–205
 - HTML coverage reports, 205–206
 - running tests, 213
 - disabling test caching, 221
 - failing fast, 220–221
 - logging in tests, 215
 - package tests, 213, 217
 - in parallel, 217–218
 - short tests, 216–217
 - specific tests, 218–219
 - with subpackages, 213–214
 - timing out, 219–220
 - verbose output, 214
- table driven testing, 206–207
 - structure of, 207–208
 - subtests, 211–213
 - writing, 208–210
- test failures
 - crafting failure messages, 201–202
 - marking, 198
 - testing.T.Error method, 198–200
 - testing.T.Fatal method, 200–201
- test file naming conventions, 195–196
- test helpers, 222
 - cleaning up, 227–229
 - defining, 222–226
 - marking functions as, 226–227
 - testing.T type, 196–197
- testing.Short function, 216–217
- testing.T type, 196–197
- testing.TB interface, 222–223
- testing.TB.Cleanup method, 227–229
- testing.TB.Helper method, 226
- testing.T.Error method, 198–200
- testing.T.Fatal method, 198, 200–201
- testing.T.Log method, 215
- testing.T.Parallel method, 217–218
- testing.T.Run method, 211
- third-party types, methods on, 178–179
- threads. *See* goroutines
- time.Ticker function, 351
- timing out
 - contexts, 407–408
 - nonresponsive shutdown, 370–371
 - tests, 219–220
- truncating files, 495–497
- type assertions, 250–251
 - avoiding panics, 284–287
 - concrete type verification, 252
 - generics and, 329–331
 - with switch statement, 252–254
 - usage example, 255–257
 - validating, 251–252
- type constraints, 315–317
 - constraints package, 326–328
 - defining, 321–322
 - mixing with method constraints, 331
 - multiple, 322–323
 - underlying, 324–326
- type definitions for arrays and slices, 83–84
- type inference, 50–51
- type keyword, 167–168
- type of value, printing, 67
- typed constants, 49
- types. *See* data types; interfaces

U

- uint numeric type, 31–32
- uint8 numeric type, 31

- uint16 numeric type, 31
- uint32 numeric type, 31
- uint64 numeric type, 31
- uintptr numeric type, 31
- unblocking channels, 343–344
- underlying type constraints, 324–326
- unidirectional channels, 351–352
- uninitialized maps, 114
- untyped constants, 50–51
- unused variables, 46–47
- Unwrapper interface, 299
- unwrapping errors, 297–301
- updating
 - complex values in maps, 125–126
 - dependencies, 21–22
- UTF-8 characters, 37–40
 - iterating over, 38–40
 - runes, 38

V

- validatable interfaces, 250
- validating type assertions, 251–252
- value receivers, pointer receivers versus, 188–189
- values
 - of arrays and slices, setting, 85
 - converting strings to/from, 72–75
 - errors as, 259–261
 - error interface, 261
 - errors.New function, 261
 - fmt.Errorf method, 262
 - handling, 262–263
 - usage example, 263–264
 - lists of. *See* lists
 - in maps
 - complex values, 123–124
 - copying on insert, 124–125
 - updating complex values, 125–126
 - panic values, capturing and returning, 269–273
 - printing
 - with formatting verbs, 67–69
 - Go-syntax of value, 69–70
 - type of value, 67
 - receiving pointers, 184
 - retrieving from contexts, 376–378
 - sending/receiving via channels, 344–345
 - wrapping contexts with, 382, 384–386
- var keyword, 40
- variables, 40–47
 - assigning, 41–42
 - functions to, 283–284
 - multiple values, 45–46
 - closures, 149–150
 - declaring, 40, 44–45
 - discarding, 47
 - exporting, 57

- functions as, 147–148
- getting pointers, 185–186
- initializing, 44–45
- modifying, 48–49
- naming, 51–57
 - bad examples, 52
 - capitalization, 57
 - package name conflicts, 53–57
 - rules for, 51–52
 - styles for, 53
- nil value, 42
- scope in if statements, 132–134
- unused, 46–47
- values
 - converting strings to/from, 72–75
 - lists of. *See* lists
 - printing, 67–69
 - printing Go-syntax of, 69–70
 - printing type of, 67
 - zero values, 42–44
- variadic arguments, 151–152
 - with append function, 89–90
 - passing slices with, 153–154
 - position of, 152–153
 - when to use, 154–156
- VCS (version control systems), modules and, 3–4
- verbose test output, 214
- verbs, formatting, 61
 - errors with, 70–71
 - escape characters, 61–63
 - explicit argument indexes, 71–72
 - floats, 66–67
 - integers, 63–66
 - strings, 63
 - value type printing, 67

W

- waitgroups
 - problem solved by, 419–421
 - sync.Waitgroup type, 421, 432–433
 - sync.WaitGroup.Add method, 422–426
 - sync.WaitGroup.Done method, 426–431
 - sync.WaitGroup.Wait method, 421–422
- walk test errors, fixing, 497–500
- walking directories, 473–477, 510–513
- Windows filepaths, 505–506
- work sharing, 339–340
- work stealing, 341
- wraparound in numeric types, 32–34
- wrapping
 - contexts with values, 382, 384–386
 - errors, 294–297
- writing
 - to closed channels, 358
 - subtests, 212–213
 - table driven testing, 208–210

Z

zero number of goroutines, adding to waitgroups, 424–425

zero values

- for arrays and slices, 81

- for closed channels, 355–357

- for maps, 121–122

- for variables, 42–44